ENGINEERING EDUCATION

Engineering Education Research and Skills for Future Development: A Comparative Modelling Study of Curriculum Configurations Under Varying Institutional Capacity

Navneet Kumar, Palak Jain, Shrashti Jaiswal, Tushar Sehrawat

Department of Information Technology, Bharati Vidyapeeth (Deemed to be University), College of Engineering, Pune, India

Published November 2026

Abstract

Engineering graduates are entering a workplace where knowing the technical fundamentals is no longer enough. Employers increasingly look for graduates who can solve unfamiliar problems, make sound design decisions, work confidently with digital tools, communicate clearly, and continue learning without constant supervision. This paper describes a controlled modelling study that compares four curriculum approaches: a lecture-centred baseline, active learning augmentation, project-based learning organised around CDIO principles, and industry-integrated project-based learning. All four are assessed using the same seven-domain competency framework, developed from accreditation criteria and published evidence on employer skill needs. To keep the comparison fair, the same simulated learner cohort is passed through every approach rather than generating a different cohort for each one. Three hundred cohorts are simulated under each of three institutional capacity conditions: well-resourced, moderately-resourced, and constrained. The study reports a composite competency attainment index, the percentage of the employer-referenced gap that is closed, a simulated delivery-cost index, ninety-five percent confidence intervals, and matched non-parametric statistical tests. Under well-resourced settings, the model gives mean attainment values of 60.11 for the baseline, 62.81 for active learning augmentation, 64.85 for project-based learning, and 65.97 for industry-integrated project-based learning, with Kendall's $W = 0.429$ and $p < 0.001$ in the Friedman test. As institutional resources become constrained, however, this ordering becomes less distinct. In the constrained scenario, the advantage of industry-integrated delivery falls to 2.70 index points, and the difference from project-based learning is no longer statistically separable after Holm adjustment at the 0.05 level. Active learning augmentation also retains much of its benefit while requiring considerably less simulated cost. The sensitivity analysis places the crossover between industry-integrated and project-based learning at about eighty-five students per staff member. Beyond that point, the most resource-intensive approach no longer produces the highest simulated attainment. These numbers come from a model input-based simulation, not from observations of an actual programme or student cohort. The aim is therefore to make the underlying reasoning transparent and repeatable, and to show when the additional effort required by high-contact teaching is supported by the model.

Keywords: Engineering education research; competency development; future skills; project-based learning; CDIO; active learning; work-integrated learning; accreditation outcomes; curriculum modelling; employability

1 Introduction

Engineering education is facing a workplace that is changing faster than many university structures can adapt. For much of the twentieth century, an undergraduate engineering programme could

reasonably focus on delivering a stable body of analytical knowledge, with the expectation that graduates would develop most of their practical skills after joining the workforce. That assumption is becoming harder to maintain. Technology cycles can be shorter than a degree programme, meaning that tools learned early in a course may already be outdated by graduation. At the same time, engineering work increasingly takes place in cross-functional teams. Graduates are expected to define loosely structured problems, work with people from outside their discipline, and pick up unfamiliar tools with relatively little guidance. These abilities can matter just as much as traditional analytical knowledge.

In response, engineering education has gradually shifted its focus from what students know to what they are able to do with that knowledge. Accreditation frameworks have been an important part of this change. The student outcomes defined by ABET and the graduate attributes developed through the Washington Accord and International Engineering Alliance focus on capabilities such as design, teamwork, communication, ethical reasoning, and continued learning [1, 2]. Indian education policy has moved in a similar direction. The National Education Policy 2020 emphasises multidisciplinary learning, flexible pathways, and experiential learning in higher education [3]. Employer studies also point toward a similar set of priorities, with analytical thinking, creative thinking, technological literacy, and confidence in learning new skills appearing prominently in skill-demand and skill-instability projections [4].

The teaching model, however, has changed more slowly. In many engineering faculties, the usual structure is still a sequence of lectures assessed mainly through individual written examinations, supported by prescribed laboratory work and a final-year project. A large body of research suggests that this approach does not develop every important outcome equally well. Freeman and colleagues, in a meta-analysis covering 225 studies in science, engineering, and mathematics, found that active learning improved examination performance by about half a standard deviation and considerably reduced failure rates compared with traditional lectures [5]. Prince reported a similar overall direction in an earlier review of engineering education [6]. Research on problem-based and project-based learning tends to show smaller effects on conventional examination results, but stronger effects when the outcome involves applying knowledge or developing practical skills [7, 8, 9].

These findings are useful, but they do not by themselves answer the curriculum decision faced by a programme director. The choice is not simply between an effect size of 0.47 and an effect size of zero. Different approaches require different levels of staffing, facilities, staff time, industry involvement, and institutional support. They are also affected differently by the conditions of a particular institution. A teaching approach that works well with twenty students per faculty member, a well-equipped makerspace, and strong industry partnerships may behave quite differently at a ratio of one hundred students per faculty member with neither the facilities nor the external network in place. Much of the existing evidence comes from relatively favourable settings, so its reported effect sizes may not transfer directly to institutions operating under tighter constraints.

Here, we address that issue by making institutional capacity part of the model rather than treating it as a fixed background condition. Four representative curriculum approaches are applied to the same simulated cohort and assessed using the same seven-domain competency framework under three levels of institutional capacity. The aim is not to introduce a new teaching method or to claim that the model measures any actual programme. Instead, the model is meant to make explicit and repeatable how the expected benefit of a high-contact approach changes when the conditions needed to deliver it become less favourable.

The work makes four main contributions. First, it defines a repeatable model in which the four curriculum approaches are tested on identical student intakes, with the model inputs and random seed stated explicitly. Second, it introduces three institutional capacity scenarios so that the dependence of each approach on resources can be examined rather than assumed away. Third, the results are presented with confidence intervals and matched non-parametric tests instead of relying only on mean values, consistent with recommendations for quantitative engineering education research [10]. Finally, attainment is considered alongside a simulated delivery-cost index. This keeps the comparison focused on both outcomes and the resources required to get them.

The rest of the paper is structured as follows. The next section reviews work on competency frameworks, instructional effectiveness, implementation conditions, and research methodology before identifying the gap addressed by this study. The subsequent sections present the objectives and research questions; describe the model, competency framework, curriculum approaches, capacity scenarios, and statistical methods; present the results and examine sensitivity to institutional capacity; discuss the findings in relation to each research question; set out the threats to validity, which are particularly important because the work is based on simulation rather than direct observation; describe the reproducibility and real-world implementation plan and outline future work; and conclude the paper.

2 Related work

2.1 Competency frameworks and the definition of the outcome

Over the past three decades, engineering education policy has increasingly moved from specifying what knowledge students should receive toward specifying the outcomes they should be able to show. ABET criteria, for example, require programmes to show that graduates can identify and solve complex engineering problems, design under realistic constraints, communicate, act ethically and professionally, work in teams, conduct experimentation, and acquire new knowledge when needed [1]. The graduate attributes developed by the International Engineering Alliance serve a similar purpose across the jurisdictions covered by the Washington Accord. Their revisions have also placed greater emphasis on sustainability, ethics, and lifelong learning [2].

The National Academy of Engineering's Engineer of 2020 reports provide another influential view of the capabilities expected from future engineers. Alongside technical depth, they emphasise adaptability, systems thinking, awareness of business and society, and leadership [11]. The Carnegie Foundation's work reached a related conclusion from the perspective of professional preparation. It argued that engineering programmes often place analytical science early in the curriculum and postpone professional practice, leaving students to experience the two as separate activities. The model so recommended introducing authentic engineering practice earlier and returning to it throughout the degree [12].

Researchers have also examined which competencies employers and practising engineers actually value. Passow and Passow's systematic review found that problem solving, communication, teamwork, data analysis, and ethics repeatedly appeared among the important competencies across different settings, although their relative importance varied by discipline and career stage [13]. Male, Bush, and Chapman reported a similar pattern in Australia, where engineers rated general capabilities such as communication and teamwork as more important to their effectiveness than several individual technical skills [14]. Employer-focused labour market studies add another consistent

theme: technological literacy, analytical ability, and the ability to keep learning are becoming increasingly important [4].

These sources provide the basis for the seven-domain competency framework used in the model and methodology section. The study does not claim that these seven domains cover every capability an engineer needs, nor does it claim that the selected weights are objectively correct. They are simulation choices intended to provide a common framework for comparing the curriculum approaches. The implications of those choices are considered again in Section VIII.

2.2 Instructional effectiveness and the evidence for active and project-based learning

One of the most widely cited quantitative findings in this area comes from Freeman and colleagues. Their meta-analysis combined 225 studies comparing active learning with traditional lecturing in undergraduate science, technology, engineering, and mathematics courses. They reported an average improvement of about 0.47 standard deviations in examination scores, along with a large reduction in failure rates [5]. The result was reasonably consistent across the disciplines and class sizes represented in their sample, and it provides the main empirical reference for the technical-depth gain assigned to active learning in the model used here.

Prince's earlier review of active learning in engineering reached a broadly similar conclusion, while also pointing out that the label covers several different practices. Collaborative learning, cooperative learning, and problem-based learning are not identical approaches, and the strength of evidence varies depending on which form and which outcome are considered [6]. Smith and colleagues provide a complementary discussion of classroom engagement practices and their basis in cooperative learning research [15]. Earlier work by Felder and Silverman also suggested that a mismatch between mostly expository teaching and the range of learning preferences within an engineering cohort could contribute to attrition [16]. The current model simplifies learner differences into a single underlying preparedness model input, so it can't capture such interactions; this limitation is discussed in Section VIII.

Earlier evidence for problem-based and project-based learning is more mixed, but it is especially relevant to the outcomes considered here. Dochy and colleagues found that problem-based learning could have a negative or negligible effect on the acquisition of declarative knowledge while improving the application of knowledge and the development of skills [8]. Strobel and van Barneveld reached a similar conclusion after synthesising several meta-analyses: problem-based learning tends to be more beneficial when outcomes involve long-term retention, applied performance, or skills, and less beneficial when the outcome is short-term recall measured through conventional examinations [7]. The model used here reflects this difference by giving project-based approaches smaller gains in technical depth than active learning, but larger gains in design, problem solving, communication, and self-directed learning.

The CDIO initiative offers a detailed framework for design-centred curriculum reform. It specifies expected learning outcomes, programme features such as design-build experiences and suitable workspaces, and a process for implementation and evaluation [17]. Graham's survey of engineering education worldwide found that institutions implementing these reforms at scale face institutional requirements that are often more demanding than the teaching methods themselves [18]. Chen, Kolmos, and Guerra reported a related finding in their review of problem-based learning implementation: the main barriers reported by institutions tend to be organisational rather than purely pedagogical [9].

2.3 Implementation context, employability and work-integrated learning

Work-integrated learning includes internships, industry-sponsored capstone projects, and cooperative education placements. Across the literature considered here, these approaches are strongly associated with graduate employability outcomes, but they also depend heavily on conditions outside the direct control of an academic department. Winberg and colleagues found consistent links between work-integrated approaches and improved graduate outcomes, while also noting that the number and quality of available placements can become the main constraint. This makes the approach difficult to scale fairly across a large student population [19].

That dependence is why institutional capacity is treated as an experimental variable in this study. An industry-integrated curriculum needs more than suitable staffing and physical facilities; it also needs reliable relationships with external partners who can provide authentic problems and supervise students. If those relationships are weak, an industry-integrated model may effectively become a project-based model using less authentic problem sources. In that situation, the additional benefit may no longer justify the additional cost. The model captures this idea through a resource sensitivity model input described in the model and methodology section, and the central question is what follows from that assumption.

2.4 Assessment, feedback and the measurement of competency

Any comparison of curriculum approaches depends on how well the chosen outcome measure represents the learning being studied, and this is a recognised challenge in engineering education research. Black and Wiliam's review highlighted the contribution of formative assessment to learning and showed how strongly conventional summative results can depend on the assessment instrument itself [20]. Hattie's synthesis of meta-analyses likewise placed feedback among the influences associated with large effects on achievement, while noting that the way feedback is implemented varies considerably [21].

The problem is especially important for outcomes such as design judgement, teamwork, and ethical reasoning because these are often the least well captured by the assessment tools that programmes already use. Litzinger and colleagues argued that professional competence develops through deliberate practice combined with useful feedback, whereas the standard undergraduate assessment system often provides too little of either [22]. Kolb's experiential learning cycle and Vygotsky's zone of proximal development are frequently used to explain why these kinds of practice environments can support learning [23, 24]. The combined attainment index used in this study deliberately abstracts away from the measurement challenge. That simplification is one of the study's most important limitations and is revisited in Section VIII.

2.5 Methodology in engineering education research

Engineering education research has developed its own expectations for how quantitative, qualitative, and mixed-methods studies should be designed and reported. Borrego, Douglas, and Amelink emphasised the importance of reporting effect sizes and uncertainty rather than relying on statistical significance alone [10]. Jamieson and Lohmann similarly argued for a systematic and scholarly approach to educational innovation, treating curriculum change as something to be studied rather than simply advocated [25]. The statistical approach in this paper follows standard procedures for matched data: the Friedman test for the overall comparison, Kendall's W as an effect-size measure of agreement, the Wilcoxon signed-rank test for pairwise checks, the Holm procedure for multiple-testing correction, and Cohen's conventions for interpreting effect magnitude [26]–[27].

2.6 Research gap

Earlier research reviewed above gives a reasonably clear picture of the competencies that matter and provides evidence that active and project-based approaches can improve many of them compared with lecture-centred delivery. It also shows that implementation depends strongly on organisational capacity. What is less clear is how the different approaches compare with one another when they are evaluated using the same outcome framework and when institutional capacity is deliberately varied. Published effect sizes generally describe the conditions under which the original studies were conducted, which are often favourable. A resource-constrained institution therefore has limited evidence for estimating how much of that benefit it can realistically achieve or whether the additional cost of a more demanding approach is justified. This study approaches that question as a transparent simulation exercise rather than as an measurement from real data.

3 Objectives and research questions

3.1 Objectives

1. Define a repeatable model of competency attainment for four representative curriculum approaches in engineering education.
2. Apply a common simulated cohort trace and a common seven-domain competency framework to all four approaches.
3. Compare well-resourced, moderately-resourced and constrained institutional capacity conditions.
4. Measure composite competency attainment, closure of the gap to an employer-based threshold, and domain-level profile.
5. Quantify differences using confidence intervals and matched non-parametric tests rather than means alone.
6. Examine attainment jointly with a simulated delivery-cost index, so that effectiveness and cost are evaluated together.
7. Determine whether the ordering of the approaches is stable as institutional capacity declines.

3.2 Research questions

RQ1: Which curriculum approach produces the highest competency attainment under well-resourced settings in the model?

RQ2: How does declining institutional capacity affect the advantage of high-contact approaches?

RQ3: How much additional attainment is gained by moving from active learning augmentation to a fully project-based or industry-integrated approach?

RQ4: What delivery cost accompanies the higher-attaining approaches in the model, and how does the cost-adjusted comparison differ from the raw ranking?

RQ5: Is the ordering of the approaches preserved as the student-to-faculty ratio increases, and if not, where does it break?

RQ6: Which competency domains are most and least sensitive to the choice of approach?

4 Model and experimental methodology

4.1 Experimental architecture

The model is organised into four connected layers. The learner-intake layer creates a simulated cohort of students, each represented by an underlying preparedness value. The curriculum-approach layer applies the gain vector associated with the selected teaching approach and reduces that gain according to the available institutional capacity. The competency-model layer combines the resulting domain scores using fixed weights to produce a combined attainment index. Finally, the assessment and measurement layer calculates the outcome measures shown for each simulated cohort. Figure 1 summarises this flow.

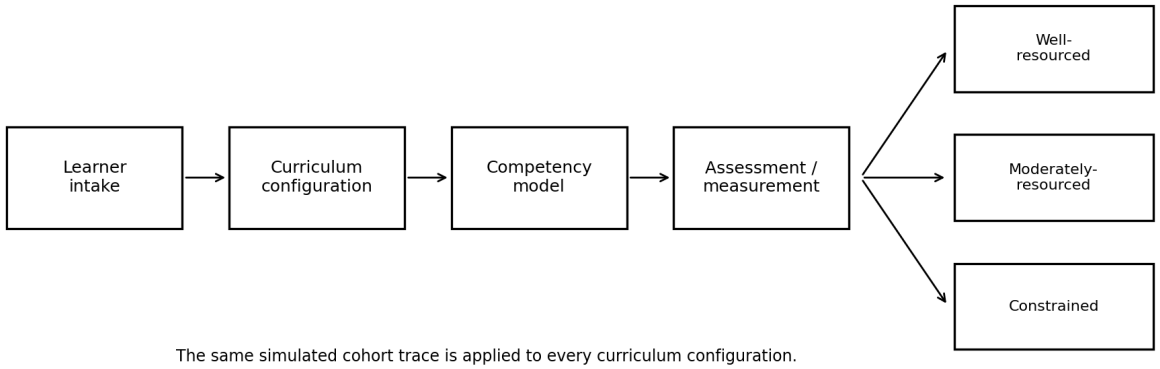


Figure 1. Model architecture. The same simulated cohort trace is applied to every configuration under every capacity scenario.

One important feature of the design is that the same cohort trace is reused. For each capacity scenario, the same 300 simulated cohorts, including the same preparedness values, are passed through all four curriculum approaches. This makes the comparison matched rather than independent. Differences caused simply by one approach receiving a more or less prepared set of students are therefore removed from the comparison, which supports the use of matched non-parametric tests. Conceptually, this is similar to keeping the student intake fixed while changing only the curriculum.

4.2 Competency framework

The simulation uses seven competency domains drawn from the overlap among ABET student outcomes, International Engineering Alliance graduate attributes, and the employer-demand literature reviewed in the related-work section [1, 4, 2, 13]. Each domain receives a fixed weight in the composite index, with the weights shown in Table 1 and summing to one. These weights are choices made for the model rather than values estimated from direct measurements. Their influence on the conclusions is therefore treated as a limitation and discussed in Section VIII.

Table 1. Competency domains and composite weights.

Domain	Weight	Principal source of the domain
Technical and disciplinary depth	0.20	ABET outcome 1; IEA attribute 1
Complex problem solving	0.18	ABET outcome 1; employer demand surveys
Design and systems thinking	0.15	ABET outcome 2; CDIO syllabus

Domain	Weight	Principal source of the domain
Digital and AI fluency	0.15	Employer demand surveys; skill-instability projections
Communication and teamwork	0.12	ABET outcomes 3 and 5; generic competency studies
Self-directed lifelong learning	0.12	ABET outcome 7; IEA attribute 12
Ethics and sustainability	0.08	ABET outcome 4; IEA attributes 8 and 9

4.3 Curriculum configurations

Four curriculum approaches are included in the model. The lecture-centred baseline represents conventional delivery: explanatory lectures, individual written examinations, prescribed laboratory sessions, and one final-year project. It serves as the reference condition, so its simulated gain is zero across all domains. Active learning augmentation keeps the basic lecture-based structure but adds techniques such as in-class activities, peer instruction, and formative assessment, drawing on practices discussed by Prince and Smith and colleagues [6, 15]. Project-based learning organised around CDIO principles changes the structure more considerably by building the curriculum around design-and-build projects throughout the degree [17]. The industry-integrated project-based approach goes one step further by adding authentic external problems, industry co-supervision, and either a placement or an industry-sponsored capstone [19].

The domain-level gains for the four approaches are shown in Table 2 in latent standard-deviation units relative to the baseline. The 0.47 technical-depth gain assigned to active learning comes directly from Freeman and colleagues' pooled examination effect [5]. The lower technical-depth gains but larger gains in design, problem solving, and communication for the project-based approaches reflect the pattern shown in the problem-based learning literature, where applied and skill outcomes tend to improve more than measures of declarative knowledge [7, 8]. The remaining values are assumptions used to build the model. They are informed by the direction and approximate size of the reviewed evidence, but they are not measurements and should not be presented as empirical effect sizes.

Table 2. Modelled domain-level gains by configuration, in latent standard deviation units, before capacity attenuation.

Domain	Baseline	Active learning	Project / CDIO	Industry-integrated
Technical and disciplinary depth	0.00	0.47	0.35	0.40
Complex problem solving	0.00	0.40	0.60	0.70
Design and systems thinking	0.00	0.20	0.80	0.85
Digital and AI fluency	0.00	0.15	0.45	0.70
Communication and teamwork	0.00	0.30	0.65	0.80
Self-directed lifelong learning	0.00	0.25	0.55	0.75
Ethics and sustainability	0.00	0.15	0.40	0.55
Weighted composite gain	0.000	0.297	0.542	0.669

4.4 Institutional capacity and resource elasticity

Each curriculum approach is given a resource-elasticity value. This represents the proportion of its simulated benefit that would be lost if institutional capacity were completely deficient. The baseline is assigned an elasticity of 0.05 because large-group lecture delivery is designed to depend relatively

little on additional capacity. Active learning receives 0.25, project-based learning 0.55, and industry-integrated project-based learning 0.75. The important simulation assumption is the ordering of these values: approaches that require more staff contact, facilities, and external partnerships are expected to lose more of their benefit when those resources become scarce. This formalises the implementation evidence reviewed in the review of implementation context [18, 19, 9].

Three institutional capacity scenarios are represented through a deficit model input. The well-resourced case has a deficit of 0.00, the moderately-resourced case 0.35, and the constrained case 0.70. For each approach, the remaining gain is calculated by multiplying its initial gain by one minus the product of its resource sensitivity and the scenario deficit. The remaining simulation settings are listed in Table 3.

Table 3. Simulation parameters.

Parameter	Value	Role
Cohorts per cell	300	Repeated random trials
Configurations	4	Compared on a shared cohort trace
Capacity scenarios	3	Well-resourced, moderate, constrained
Competency domains	7	Weighted composite outcome
Baseline index mean	60.0	Reference attainment level
Index points per latent SD	9.0	Scale mapping
Employer-referenced threshold	80.0	Target for gap reduction metric
Cohort variance (SD)	0.45	Between-cohort preparedness
Residual variance (SD)	0.30	Within-cell random term
Capacity deficit levels	0.00, 0.35, 0.70	Scenario definition
Delivery cost index	1.00, 1.12, 1.48, 1.85	Modelled relative cost per student
Random seed	20260916	Reproducibility

4.5 Outcome metrics

For every simulated cohort, the model calculates three main outcomes. The composite competency attainment index starts from a baseline mean of 60 and adds nine index points for each latent standard deviation represented by the combination of cohort preparedness, the retained weighted gain of the curriculum approach, and a residual random term. Values are clipped to the 0–100 range. The gap-closure measure expresses the increase above the baseline as a percentage of the distance between the baseline and the employer-based threshold of 80. Put simply, it shows how much of the simulated shortfall is removed. The delivery-cost index is a fixed multiplier for each approach and represents the simulated relative cost per student, including staff contact time, facilities, and partnership management. It is a model input, not a real-world observation, and is shown so that attainment is considered together with cost.

It is worth being clear about what the composite index does and does not represent. It is a construct created for the model; it is not an examination score, assessment rubric, or validated measurement instrument. Combining seven different competency domains into a single number assumes that they can be placed on a common scale and weighted linearly, which is a strong and debatable assumption.

The index is used because it makes the comparison transparent and repeatable, not because the study claims that it is a complete measurement model for engineering competence.

4.6 Statistical analysis

Because all four approaches use the same cohort trace, their attainment scores are matched. The overall comparison is therefore performed with the Friedman test [28], while Kendall's W is shown as an effect-size measure of concordance among the approach rankings [27]. Pairwise differences are tested with the Wilcoxon signed-rank test [29], and the p-values from the six pairwise checks within each scenario are adjusted using the Holm sequentially rejective procedure [30]. A two-tailed significance level of 0.05 is used throughout. Ninety-five percent confidence intervals for the means are calculated using a normal approximation, and effect magnitudes are interpreted using Cohen's conventions [26].

4.7 Reproducibility

The model is implemented in Python, using NumPy for random generation, SciPy for the statistical tests, and Matplotlib for the figures. The random seed is 20260916, and 300 cohorts are simulated for every approach and capacity scenario. All other model inputs are those listed in Tables 1, 2 and 3. Cohort generation is kept separate from approach evaluation, so an additional curriculum approach could be added by specifying its gain vector and resource sensitivity without changing the learner-intake or statistical procedures.

5 Results

The results in this section come entirely from the model described in the model and methodology section. They are not observations of competency attainment from an actual university programme, and the numerical values should not be interpreted as predictions of what a particular institution will achieve. The purpose of the comparison is narrower: to show the direction and approximate size of the differences produced by the stated assumptions, and to examine whether the relative ordering of the approaches remains stable as institutional capacity changes.

Table 4 presents the competency attainment, gap closure and delivery cost.

Table 4. Competency attainment, gap closure and delivery cost.

Scenario	Configuration	Attainment index	95% CI	Gap closure (%)	Cost index
Well-resourced	Lecture-centred baseline	60.113	59.539 to 60.686	0.56	1.00
Well-resourced	Active learning augmentation	62.812	62.226 to 63.398	14.06	1.12
Well-resourced	Project-based / CDIO	64.848	64.274 to 65.422	24.24	1.48
Well-resourced	Industry-integrated PBL	65.968	65.347 to 66.589	29.84	1.85
Moderate	Lecture-centred baseline	60.113	59.539 to 60.686	0.56	1.00

Scenario	Configuration	Attainment index	95% CI	Gap closure (%)	Cost index
Moderate	Active learning augmentation	62.579	61.992 to 63.165	12.89	1.12
Moderate	Project-based / CDIO	63.910	63.336 to 64.483	19.55	1.48
Moderate	Industry-integrated PBL	64.389	63.767 to 65.010	21.94	1.85
Constrained	Lecture-centred baseline	60.113	59.539 to 60.686	0.56	1.00
Constrained	Active learning augmentation	62.345	61.759 to 62.931	11.73	1.12
Constrained	Project-based / CDIO	62.972	62.398 to 63.545	14.86	1.48
Constrained	Industry-integrated PBL	62.809	62.188 to 63.431	14.05	1.85

Figure 2 compares the mean attainment index of the four configurations across the three capacity scenarios.

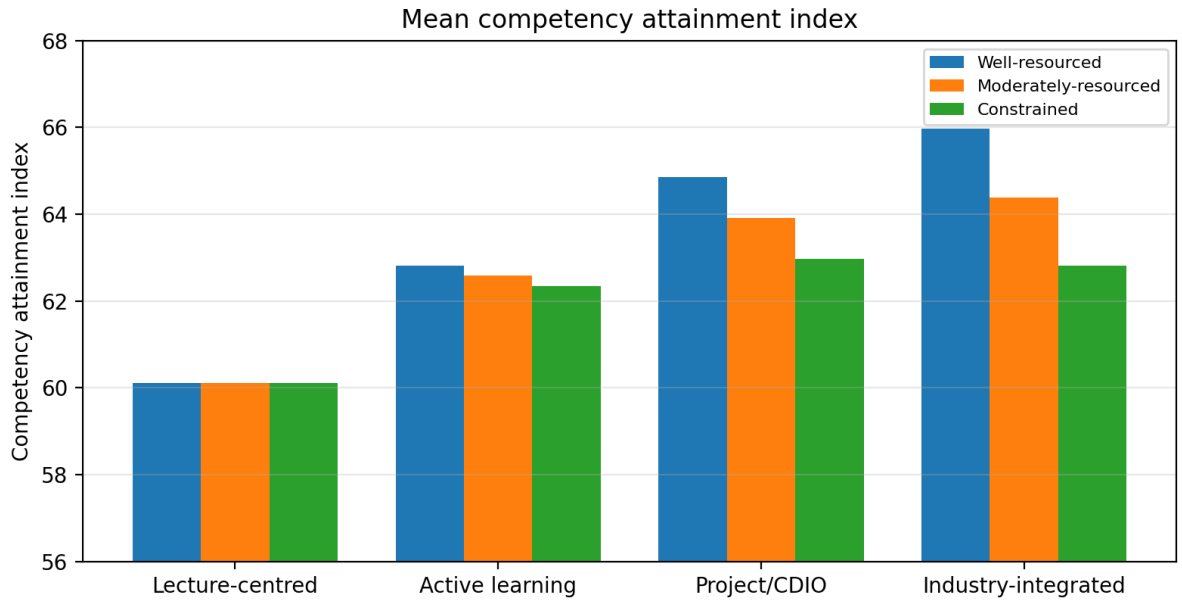


Figure 2. Mean competency attainment index by configuration across the three institutional capacity scenarios.

5.1 Well-resourced conditions

The well-resourced scenario gives the clearest separation among the four approaches. The lecture-centred baseline has a mean attainment index of 60.11, with a ninety-five percent confidence interval of 59.54–60.69. Active learning augmentation reaches 62.81, project-based learning reaches 64.85, and industry-integrated project-based learning reaches 65.97. The corresponding confidence intervals do not overlap. This is also the condition in which the model shows the clearest benefit for approaches that involve more direct student interaction and practical work.

When the results are expressed as closure of the gap to the employer-based threshold of 80, the baseline closes only 0.56% of the simulated shortfall. Active learning augmentation closes 14.06%, project-based learning 24.24%, and industry-integrated project-based learning 29.84%. An important point is the ceiling implied by these values: even under the most favourable conditions represented in the model, the strongest approach closes less than one-third of the gap. Given the gain assumptions used in the simulation, this suggests that curriculum structure by itself is unlikely to account for the entire competency shortfall described in employer studies.

5.2 Moderately-resourced conditions

With a moderate resource deficit of 0.35, the same basic ordering remains, but the differences become noticeably smaller. The baseline stays at 60.11 because its low elasticity makes it almost insensitive to the capacity change by design. Active learning falls only slightly, from 62.81 to 62.58, a reduction of 0.23 index points. Project-based learning falls from 64.85 to 63.91, while industry-integrated project-based learning drops from 65.97 to 64.39. The more resource-dependent approaches so lose more of their simulated advantage as capacity declines.

As a result, the gap between the two project-based approaches shrinks from 1.12 index points in the well-resourced case to 0.48 points in the moderate case. The Holm-adjusted Wilcoxon comparison between them changes from a p-value of 2.46×10^{-6} in the well-resourced scenario to 0.0448 in the moderate scenario. The latter remains below 0.05, but only narrowly. At the same time, industry integration carries an additional 0.37 on the cost index compared with project-based learning. Within the assumptions of the model, the extra cost is therefore associated with a difference that is becoming harder to distinguish statistically.

5.3 Constrained conditions

The constrained scenario changes the picture more clearly. Industry-integrated project-based learning reaches 62.81, slightly below the 62.97 achieved by project-based learning, even though its simulated cost is twenty-five percent higher. The Holm-adjusted Wilcoxon test gives $p = 0.434$ for the comparison between the two, so the model does not distinguish them at the 0.05 level. Put simply, under these particular assumptions, the additional industry integration no longer produces a statistically separable attainment benefit over the cheaper project-based approach.

Active learning augmentation reaches 62.35 under the constrained condition, losing only 0.47 index points compared with its well-resourced value of 62.81. It is within 0.62 points of project-based learning and 0.46 points of industry-integrated project-based learning, while its cost index remains 1.12 compared with 1.48 and 1.85. The same pattern appears in the gap-closure figures: active learning closes 11.73% of the simulated gap, compared with 14.86% for project-based learning and 14.05% for industry-integrated learning, despite the latter two carrying considerably higher cost indices.

Figure 3 shows the corresponding closure of the competency gap.

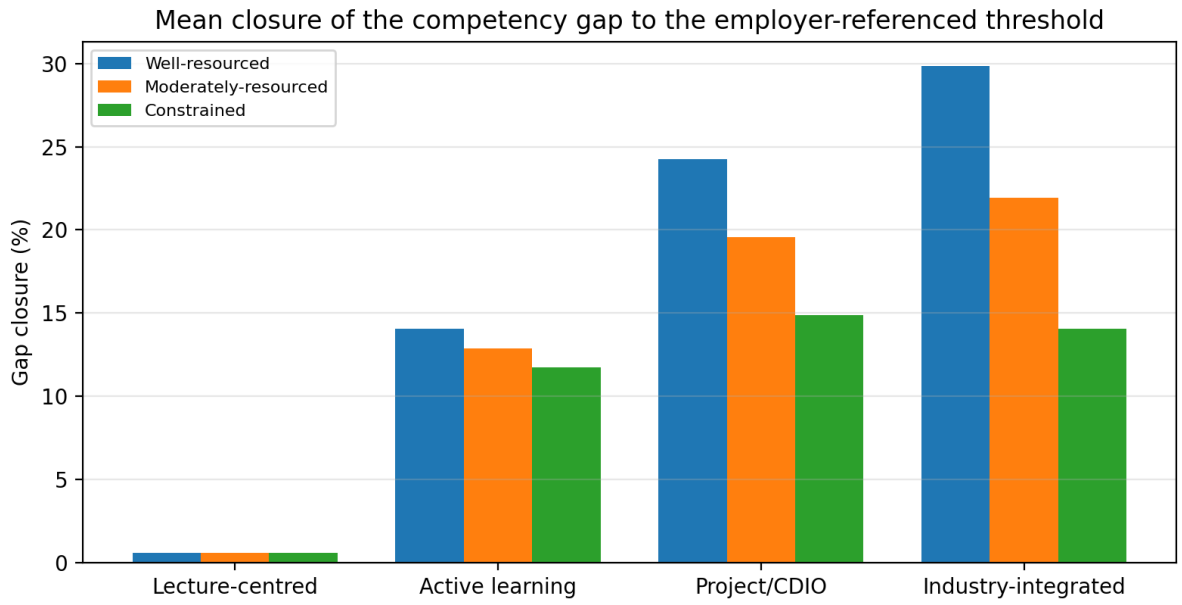


Figure 3. Mean closure of the competency gap to the employer-referenced threshold, by configuration and capacity scenario.

5.4 Statistical comparison

The Friedman test finds statistically detectable differences among the four approaches in all three scenarios, but the size of the overall effect becomes smaller as institutional capacity deteriorates. In the well-resourced case, the Friedman chi-square is 385.66 with $p < 0.001$ and Kendall's $W = 0.429$. For the moderately-resourced case, the statistic is 266.48 with $p < 0.001$ and $W = 0.296$. Under resource-constrained settings, it falls to 140.57 with $p < 0.001$ and $W = 0.156$.

The fall in Kendall's W from 0.429 to 0.156 is one of the clearest statistical signals in the study. With 300 matched cohorts in each cell, even fairly small differences can reach statistical significance, so the p-values alone say little about practical importance. What changes more meaningfully is the degree of agreement in the approach ordering. As capacity deteriorates, the approaches become less clearly separated, and the choice of curriculum accounts for less of the variation represented by the model. The small effect size in the constrained scenario also suggests that statistical significance by itself should not be treated as a reason for a major reallocation of institutional resources.

Table 5 presents the global statistical test by scenario.

Table 5. Global statistical test by scenario.

Scenario	Friedman chi-square	p value	Kendall W	Interpretation
Well-resourced	385.66	< 0.001	0.429	Moderate concordance
Moderately-resourced	266.48	< 0.001	0.296	Small to moderate
Constrained	140.57	< 0.001	0.156	Small

The Holm-adjusted pairwise checks show the same general pattern. Every comparison between the baseline and a non-baseline approach is significant in all three scenarios. The comparisons among the three more active approaches are more sensitive to capacity. More specifically, the project-based versus industry-integrated comparison changes from a very small p-value in the well-resourced

scenario to 0.0448 in the moderate scenario and then to 0.434 in the constrained scenario. This comparison captures much of the practical change produced by the capacity assumptions.

Table 6 presents the selected Holm-adjusted pairwise comparisons.

Table 6. Selected Holm-adjusted pairwise comparisons.

Comparison	Well-resourced	Moderate	Constrained
Baseline vs active learning	< 0.001	< 0.001	< 0.001
Baseline vs project-based	< 0.001	< 0.001	< 0.001
Baseline vs industry-integrated	5.4×10^{-47}	7.3×10^{-39}	< 0.001
Active learning vs project-based	< 0.001	< 0.001	< 0.001
Project-based vs industry-integrated	2.46×10^{-6}	0.0448	0.434

5.5 Domain-level profile

Looking only at the composite index can hide important differences between competency domains. The smallest spread among the approaches occurs in technical and disciplinary depth, where the initial gains range from 0.35 to 0.47 standard deviations and active learning has the largest value. The biggest differences appear in design and systems thinking, where the gains range from 0.20 to 0.85, and in self-directed lifelong learning, where they range from 0.25 to 0.75. Digital and AI fluency also shows a notable separation between the two project-based approaches, with gains of 0.45 and 0.70 respectively. In the model, this reflects the assumption that authentic industry problems expose students to current tools and workflows that internally defined projects may not always provide.

These results suggest that the four approaches should not be viewed as interchangeable versions of the same teaching model. If the main concern is examination performance in the disciplinary core, the model gives active learning a large share of the available gain at a cost index of 1.12. For outcomes such as design judgement, professional communication, and independent learning, the project-based approaches have much larger assumed gains. This is why the study treats the competency domains separately rather than reducing the curriculum choice to a single general measure of quality.

6 Sensitivity to institutional capacity

The three capacity scenarios are complemented by a continuous sensitivity analysis based on the student-to-faculty ratio. The ratio is varied from 20:1, representing a well-resourced programme, to 120:1, representing a severely constrained programme. The resource deficit is calculated from the amount by which the ratio exceeds 20 and is capped at one. At each ratio, 300 cohorts are simulated. The resulting mean attainment values are shown in Table 7 and visualised in Figure 4.

Table 7. Sensitivity of mean attainment to student-to-faculty ratio.

Ratio	Baseline	Active learning	Project / CDIO	Industry-integrated
20 : 1	59.77	62.61	64.85	65.65
40 : 1	60.09	62.46	64.22	64.77
60 : 1	59.79	62.13	63.43	64.20
80 : 1	59.55	61.73	62.78	62.95

Ratio	Baseline	Active learning	Project / CDIO	Industry-integrated
100 : 1	59.68	61.62	62.24	61.67
120 : 1	59.95	61.32	61.81	61.24

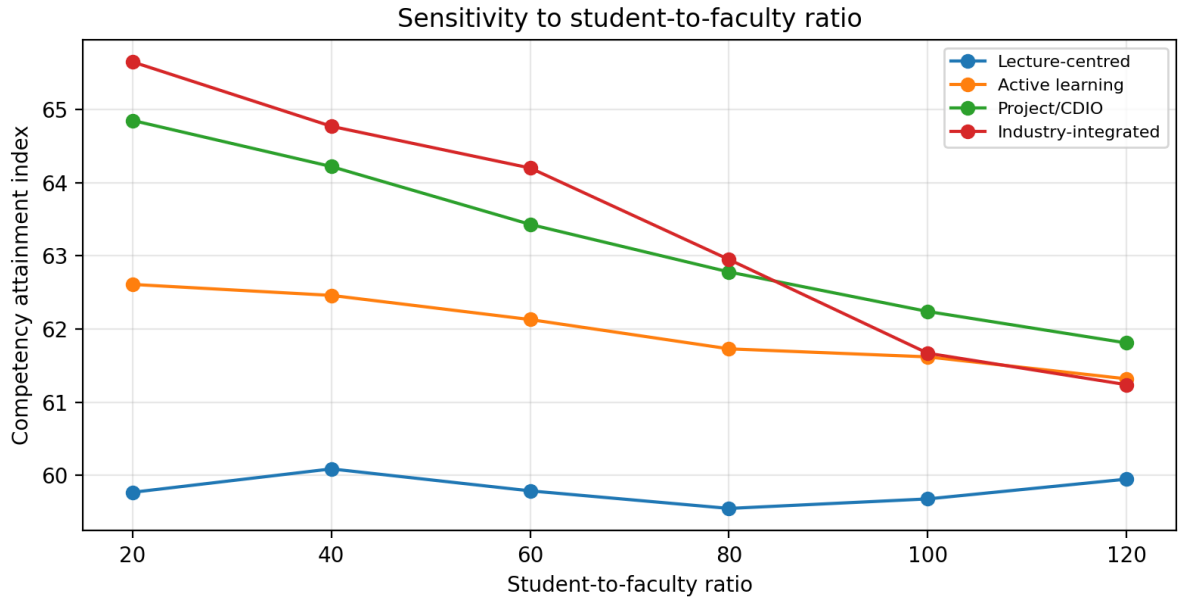


Figure 4. Mean competency attainment index across student-to-faculty ratios from 20:1 to 120:1.

Figure 4 highlights three patterns. First, the baseline remains almost flat across the entire range, varying between 59.55 and 60.09. This small variation is random rather than a systematic response to the student-to-faculty ratio. In the model, a lecture-centred programme therefore produces almost the same attainment at 120 students per staff member as it does at 20. That stability comes from the baseline's low resource sensitivity, but it also means that the simulated baseline closes very little of the competency gap.

Second, the approaches that depend more heavily on resources decline more sharply as the student-to-faculty ratio increases. Industry-integrated project-based learning falls from 65.65 at 20:1 to 61.24 at 120:1, a drop of 4.41 index points. Project-based learning falls from 64.85 to 61.81, a reduction of 3.04 points. Active learning augmentation declines from 62.61 to 61.32, a smaller reduction of 1.29 points. The order of these declines follows the resource elasticities built into the model; the more interesting question is where the resulting curves cross.

The third pattern is the crossover between industry-integrated and project-based learning. Their curves meet at about 85 students per staff member. Below that ratio, the industry-integrated approach has the higher simulated attainment, while above it the project-based approach becomes higher. At 120:1, industry-integrated learning falls to 61.24, below project-based learning at 61.81 and active learning at 61.32, while still carrying the highest cost index. At the end of the sweep, the three non-baseline approaches are separated by only 0.57 index points, so their differences in attainment become quite small within the model.

The crossover around 85:1 should not be treated as an externally validated threshold. It comes directly from the elasticity values assumed in the subsection on institutional capacity and resource elasticity, so changing those assumptions would move the point of intersection. What is more general

is the possibility of a crossover itself. If a more expensive approach is also more dependent on institutional capacity than a cheaper alternative, there can be a point at which declining capacity removes its additional benefit. The sensitivity analysis is meant to show that mechanism rather than establish a universal student-to-faculty threshold.

7 Discussion

7.1 Answer to RQ1

Under well-resourced settings, the model produces its highest mean attainment for industry-integrated project-based learning, at 65.97, followed by project-based learning at 64.85, active learning augmentation at 62.81, and the lecture-centred baseline at 60.11. The confidence intervals do not overlap, and the Friedman test gives Kendall's W = 0.429. This pattern is broadly consistent with the direction suggested by the literature considered earlier. In this study, however, its main purpose is to provide a reference point for examining how the differences change when institutional capacity is reduced.

7.2 Answer to RQ2

Reducing institutional capacity does not affect all approaches equally. The model reduces each approach's benefit according to its assumed dependence on staff time, facilities, and external partnerships. Between the well-resourced and constrained scenarios, active learning loses 0.47 index points, project-based learning loses 1.88, and industry-integrated project-based learning loses 3.16. As a result, the upper part of the ordering changes: under constraint, industry-integrated delivery reaches 62.81, slightly below project-based learning at 62.97, and the two are statistically indistinguishable with a Holm-adjusted p-value of 0.434.

This behaviour is consistent with a recurring theme in the implementation literature. Graham's survey of institutions undertaking large-scale curriculum reform and the scoping review by Chen and colleagues both describe institutional barriers as important obstacles to reproducing reforms across different settings [18, 9]. The model used here provides one possible explanation for that pattern: approaches that require the most favourable delivery conditions are also the approaches whose benefits are most sensitive to those conditions. The result should so be understood as an implication of the model rather than as direct evidence from observed institutions.

7.3 Answer to RQ3

The additional attainment gained by moving beyond active learning depends strongly on the capacity scenario. In the well-resourced case, moving from active learning augmentation to project-based learning adds 2.04 index points, and moving from project-based to industry-integrated learning adds another 1.12 points. Under resource-constrained settings, the corresponding changes are only 0.63 points and −0.16 points. Thus, the second step actually reduces the simulated attainment under the constrained assumptions. In terms of gap reduction, the full move from active learning to industry-integrated delivery adds 15.78 percentage points when resources are available but only 2.32 points under constraint.

For an institution considering such a change, the model suggests that the local delivery conditions matter at least as much as the general teaching approach. The simulated return from industry integration is concentrated in settings that can support the staffing, facilities, and external

relationships it requires. Where those conditions are not available, active learning captures a large portion of the available benefit at a much smaller incremental cost: 0.12 on the cost index compared with 0.85 for moving from active learning to industry-integrated delivery.

7.4 Answer to RQ4

The model assigns cost indices of 1.00, 1.12, 1.48, and 1.85 to the four approaches. Dividing gap reduction by cost gives a simple, illustrative efficiency measure. Under well-resourced settings, the values are 0.56, 12.55, 16.38, and 16.13 for the baseline, active learning, project-based, and industry-integrated approaches respectively. Under resource-constrained settings, they become 0.56, 10.47, 10.04, and 7.59. These values show how including cost can change the interpretation of the same attainment results.

The change between these two sets of efficiency figures is one of the main practical points of the paper. Neither set of figures is intrinsically correct or incorrect; they describe different institutional conditions. If curriculum effectiveness is shown without considering delivery cost, and if most of the supporting evidence comes from well-resourced institutions, the resulting picture may not transfer to institutions with much tighter constraints. The model therefore supports reporting attainment and delivery conditions together rather than presenting effectiveness as a context-free property of a curriculum approach.

7.5 Answer to RQ5

The ordering does not remain unchanged across the full sensitivity range. The industry-integrated and project-based curves cross at roughly 85 students per staff member, and by 120:1 the three non-baseline approaches are within 0.57 index points of one another. The exact crossover point is a consequence of the elasticity values assumed in the model and so has no generalisability beyond the simulation on its own. The broader pattern, however, follows from the model's assumption that the more expensive approach is also more dependent on institutional capacity, an assumption that is supported qualitatively by the implementation literature [19, 9].

7.6 Answer to RQ6

The effect of curriculum approach is not equally strong across all competency domains. Technical and disciplinary depth is the least sensitive, with initial gains ranging from 0.35 to 0.47 standard deviations and active learning showing the largest gain. Design and systems thinking is the most sensitive, ranging from 0.20 to 0.85. Self-directed lifelong learning and communication and teamwork also show large differences. This pattern is consistent with the literature suggesting that problem-based approaches may be less effective for short-term declarative knowledge while being more useful for application and skill development [7, 8]. It also reinforces the idea that a curriculum should be considered in relation to the particular competency profile it is meant to develop.

7.7 Relation to prior work

The overall direction of the model results is consistent with the literature considered in the related-work section. Active and project-based approaches show higher simulated attainment than lecture-centred delivery, their largest assumed gains occur in skill-oriented outcomes rather than purely disciplinary examination performance, and work-integrated approaches show their greatest benefit when the necessary conditions can be sustained. The distinctive contribution here is the conditional structure. By making institutional capacity an explicit variable and assigning each approach its own

resource sensitivity, the model allows the relative ordering to change instead of assuming that one ordering applies everywhere. To the authors' knowledge, this conditional relationship is more often discussed informally in implementation studies than represented explicitly in a model.

8 Threats to validity

The limitations of this study are large, and they need to be made explicit because the results could otherwise be mistaken for empirical findings. This is a modelling study, so its numbers describe what follows from the assumptions built into the model rather than what has been directly observed in universities.

The most important limitation is straightforward: no students were actually observed. There are no measurements from a student, programme, or institution. Every number shown in the results and sensitivity sections is generated from the model inputs defined in the model and methodology section. The model shows what those assumptions imply, but it does not establish that the same values would occur in the real-world setting. For instance, a statement such as “project-based learning produces an attainment index of 64.85” refers only to the simulated outcome and should be interpreted in that context.

The gain model inputs are only loosely connected to published effect sizes. The 0.47 value used for technical depth under active learning comes from a genuine meta-analytic result [5], but that result concerns examination scores across a pooled sample, not a seven-domain competency measure of the type used here. The conversion from that evidence to the model's domain gain is therefore an assumption. The other gain values, especially those assigned to industry-integrated learning, reflect the direction and approximate magnitude of the literature considered but are not taken from a single measurement from real data. They should not be presented as measured effect sizes.

The resource-elasticity model inputs are the least firmly grounded part of the model, yet they drive much of its central result. For instance, assigning an elasticity of 0.75 to industry-integrated learning and 0.25 to active learning means that the former loses much more of its simulated advantage when capacity is deficient. This ordering is supported qualitatively by the implementation literature, but the exact magnitudes are choices made for the model rather than measurements. As a result, the existence of a crossover is more informative than its exact location. A different set of elasticity values would move the crossover point.

The combined attainment index also makes a strong simplifying assumption: seven different competency domains can be combined on one scale using fixed linear weights. That assumption can be questioned for several reasons. The domains are related to one another; for example, design ability and problem solving can interact. The selected weights are choices made for the model, and this work does not report a sensitivity analysis for alternative weights. Reducing complex outcomes such as ethical reasoning or teamwork to a single scalar value that can be averaged with technical depth is also a large abstraction from the underlying measurement problem [22, 20].

Institutional capacity is represented by only one dimension in the model. In practice, staffing, physical facilities, industry partnerships, staff development, and institutional incentives can vary independently and can affect different curriculum approaches in different ways. An institution might have strong facilities but weak industry connections, or strong external partnerships but insufficient staff time. The single deficit model input used here cannot capture these combinations.

Other potentially important factors are also omitted. Student motivation and its interaction with the teaching approach are not simulated. Faculty capability and the effect of professional development on how well a approach is implemented are also absent, even though the implementation literature identifies them as important. Selection effects are not included either; institutions that adopt industry-integrated approaches may differ systematically from those that do not. The model also ignores transition costs, meaning that it compares steady-state approaches rather than the disruption involved in moving from one system to another. Finally, the equity of access to placements, which Winberg and colleagues identify as an important constraint on work-integrated learning [19], is not represented.

Statistical significance should also be interpreted cautiously. With 300 matched cohorts in each cell, even a small fraction of an index point can become statistically detectable. This is why the Friedman test is significant in all three scenarios, including the case where Kendall's W is small. The p-values mainly describe the precision of the simulation under the specified assumptions; they do not measure practical importance. The effect sizes and confidence intervals are more informative for that purpose, and the decline in Kendall's W from 0.429 to 0.156 captures more of the substantive change than any individual p-value.

9 Reproducibility and empirical implementation plan

The model can be reproduced directly from the information given in the model and methodology section because the model inputs, weights, elasticities, and random seed are all specified. Cohort generation is separated from approach evaluation, so adding another curriculum approach would require only a new gain vector and resource-elasticity value. This separation also makes it easier to modify or extend the model without changing the underlying student-generation process.

The natural next step is to test the model using real data. A multi-institution study could include programmes using each of the four approaches across different student-to-faculty ratios. Rather than relying on one composite score, the seven competency domains could be measured using instruments suited to each domain: validated concept inventories for technical depth, rubric-based assessment of design artefacts for design and systems thinking, structured observation of team processes for communication and teamwork, and self-regulated learning measures for the self-directed learning domain. Domain-level results should be shown separately, with any composite measure presented as an additional summary rather than a replacement for the individual measures.

Such a study would allow the resource elasticities to be estimated from data instead of assumed. These are the model inputs about which the current model has the least empirical confidence and on which its central crossover result depends most strongly. Incoming student preparedness would also need to be measured and controlled, since the matched cohort design used in the simulation cannot be reproduced perfectly in a real study. Delivery costs should be recorded using observable quantities such as staff contact hours, facility use, and partnership-management effort rather than a fixed index. With these measurements, the conditional structure proposed by the model could be tested directly, including whether a crossover appears at all and, if it does, where it occurs.

Algorithm 1. Simulation procedure

1. Set the random seed and load the approach model inputs from Tables 1, 2 and 3.
2. Create one cohort trace containing 300 underlying preparedness values for the scenario.
3. Run that same cohort trace through each of the four curriculum approaches.

4. Reduce each approach's domain-level gain according to its resource sensitivity and the resource deficit.
5. Combine the adjusted domain gains using the fixed competency weights to get the composite gain.
6. Calculate the attainment index, gap reduction, and delivery-cost index for each simulated cohort.
7. Repeat the process for the three capacity scenarios and the six student-to-faculty ratios.
8. Calculate the means and ninety-five percent confidence intervals, then run the Friedman test with Kendall's W and the Holm-adjusted Wilcoxon signed-rank comparisons on the matched samples.

10 Future work

- Test the model against measured competency outcomes in a multi-institution study covering different student-to-faculty ratios.
- Estimate the resource elasticities from observed data rather than treating them as assumed model inputs, since they strongly influence the main result.
- Replace the combined attainment index with separate, instrument-specific measures for each competency domain.
- Expand the one-dimensional capacity model into separate components for staffing, facilities, industry partnerships, and staff development.
- Include faculty capability and the effect of staff development on how successfully each approach is implemented.
- Represent transition costs and the multi-year process of moving between curriculum approaches instead of comparing only steady-state conditions.
- Include equity of access to placements and examine the distributional effects of work-integrated learning opportunities.
- Add hybrid approaches, such as a project spine combined with a lecture-centred core, which are not represented by the current four approaches.
- Extend the competency framework to treat artificial intelligence literacy as a separate domain as AI becomes more integrated into engineering practice.
- Test how sensitive the results are to alternative choices of the composite competency weights.

11 Conclusion

This paper describes a controlled modelling study of four engineering-education curriculum approaches, evaluated against seven competency domains under three levels of institutional capacity. Three hundred simulated cohorts are used for each cell, and the same cohort trace is passed through every approach so that the comparisons remain matched.

In the well-resourced scenario, the model produces the ordering expected from the literature, with industry-integrated project-based learning at 65.97, project-based learning at 64.85, active learning augmentation at 62.81, and the lecture-centred baseline at 60.11. Kendall's W is 0.429. Even the highest-attaining approach closes less than one-third of the simulated gap to the employer-based threshold. This reinforces the point that, within the assumptions of the model, changing curriculum structure alone does not remove the full competency gap described in employer studies.

The ordering changes when institutional capacity is constrained. The four approaches depend on staff contact, facilities, and external partnerships to different degrees, so their simulated benefits decline at different rates. Under the constrained scenario, industry-integrated project-based learning falls below project-based learning and the difference between them is not statistically separable at a Holm-adjusted p of 0.434. At the same time, industry-integrated learning still carries a twenty-five percent higher simulated cost, and Kendall's W falls from 0.429 to 0.156. The sensitivity analysis places the simulated crossover near 85:1 in the student-to-faculty ratio, although that exact point should not be treated as externally valid.

Two broader reporting lessons follow from the analysis, rather than a recommendation to adopt one particular pedagogy. First, curriculum interventions should be considered using both attainment and delivery cost, because the efficiency picture changes between favourable and resource-constrained settings even when the raw attainment differences remain relatively similar. Second, the conditions under which a curriculum is delivered should be treated as a study variable rather than background information. An effect size observed at a 20:1 student-to-faculty ratio does not automatically describe a programme operating at 100:1, and the model shows that the difference can be large enough to change the relative ordering of approaches.

The findings should ultimately be read as the output of a parameter-based model, rather than as measurements from actual programmes. The resource elasticities that drive the main result are also the least supported by observed data in the model. Measuring those relationships in actual institutions is therefore the most important next step if the conditional structure proposed here is to be tested.

Acknowledgement

The authors gratefully acknowledge the Department of Information Technology, Bharati Vidyapeeth (Deemed to be University), College of Engineering, Pune, India, for providing the academic environment and support needed to carry out this work. The authors also appreciate the guidance and encouragement received during the development of this modelling study on curriculum approaches and competency development in engineering education. Finally, they thank everyone who contributed, directly or indirectly, to completing the work.

References

[1] ABET Engineering Accreditation Commission. *Criteria for accrediting engineering programs*. Baltimore, MD: ABET; 2023.

[2] International Engineering Alliance. *Graduate attributes and professional competencies, version 4*. Wellington: International Engineering Alliance; 2021.

[3] Ministry of Education, Government of India. *National Education Policy 2020*. New Delhi: Government of India; 2020.

[4] World Economic Forum. *The future of jobs report 2023*. Geneva: World Economic Forum; 2023.

[5] Freeman S, Eddy SL, McDonough M, Smith MK, Okoroafor N, Jordt H, et al. Active learning increases student performance in science, engineering, and mathematics. *Proceedings of the National Academy of Sciences*. 2014;111(23):8410–8415. https://doi.org/10.1073/pnas.1319030111

[6] Prince M. Does active learning work? A review of the research. *Journal of Engineering Education*. 2004;93(3):223–231. https://doi.org/10.1002/j.2168-9830.2004.tb00809.x

[7] Strobel J, van Barneveld A. When is PBL more effective? A meta-synthesis of meta-analyses comparing PBL to conventional classrooms. *Interdisciplinary Journal of Problem-Based Learning*. 2009;3(1):44–58. https://doi.org/10.7771/1541-5015.1046

[8] Dochy F, Segers M, Van den Bossche P, Gijbels D. Effects of problem-based learning: A meta-analysis. *Learning and Instruction*. 2003;13(5):533–568. https://doi.org/10.1016/S0959-4752(02)00025-7

[9] Chen J, Kolmos A, Guerra A. Implementation and challenges of problem-based learning in higher education: A scoping review. *European Journal of Engineering Education*. 2021;46(1):90–115. https://doi.org/10.1080/03043797.2020.1718615

[10] Borrego M, Douglas EP, Amelink CT. Quantitative, qualitative, and mixed research methods in engineering education. *Journal of Engineering Education*. 2009;98(1):53–66. https://doi.org/10.1002/j.2168-9830.2009.tb01005.x

[11] National Academy of Engineering. *The engineer of 2020: Visions of engineering in the new century*. Washington, DC: National Academies Press; 2004.

[12] Sheppard SD, Macatangay K, Colby A, Sullivan WM. *Educating engineers: Designing for the future of the field*. San Francisco, CA: Jossey-Bass; 2008.

[13] Passow HJ, Passow CH. What competencies should undergraduate engineering programs emphasize? A systematic review. *Journal of Engineering Education*. 2017;106(3):475–526. https://doi.org/10.1002/jee.20171

[14] Male SA, Bush MB, Chapman ES. An Australian study of generic competencies required by engineers. *European Journal of Engineering Education*. 2011;36(2):151–163. https://doi.org/10.1080/03043797.2011.569703

[15] Smith KA, Sheppard SD, Johnson DW, Johnson RT. Pedagogies of engagement: Classroom-based practices. *Journal of Engineering Education*. 2005;94(1):87–101. https://doi.org/10.1002/j.2168-9830.2005.tb00831.x

[16] Felder RM, Silverman LK. Learning and teaching styles in engineering education. *Engineering Education*. 1988;78(7):674–681.

[17] Crawley EF, Malmqvist J, Östlund S, Brodeur DR, Edström K. *Rethinking engineering education: The CDIO approach*. 2nd ed. Cham: Springer; 2014.

[18] Graham R. *The global state of the art in engineering education*. Cambridge, MA: Massachusetts Institute of Technology; 2018.

[19] Winberg C, Bramhall M, Greenfield D, Johnson P, Rowlett P, Lewis O, et al. Developing employability in engineering education: A systematic review of the literature. *European Journal of Engineering Education*. 2020;45(2):165–180. https://doi.org/10.1080/03043797.2018.1534086

[20] Black P, Wiliam D. Assessment and classroom learning. *Assessment in Education: Principles, Policy and Practice*. 1998;5(1):7–74. https://doi.org/10.1080/0969595980050102

[21] Hattie J. *Visible learning: A synthesis of over 800 meta-analyses relating to achievement*. London: Routledge; 2009.

[22] Litzinger TA, Lattuca LR, Hadgraft RG, Newstetter WC. Engineering education and the development of expertise. *Journal of Engineering Education*. 2011;100(1):123–150. https://doi.org/10.1002/j.2168-9830.2011.tb00006.x

[23] Kolb DA. *Experiential learning: Experience as the source of learning and development*. Englewood Cliffs, NJ: Prentice Hall; 1984.

[24] Vygotsky LS. *Mind in society: The development of higher psychological processes*. Cambridge, MA: Harvard University Press; 1978.

[25] Jamieson LH, Lohmann JR. *Creating a culture for scholarly and systematic innovation in engineering education*. Washington, DC: American Society for Engineering Education; 2009.

[26] Cohen J. *Statistical power analysis for the behavioral sciences*. 2nd ed. Hillsdale, NJ: Lawrence Erlbaum Associates; 1988.

[27] Kendall MG, Babington Smith B. The problem of m rankings. *Annals of Mathematical Statistics*. 1939;10(3):275–287. https://doi.org/10.1214/aoms/1177732186

[28] Friedman M. The use of ranks to avoid the assumption of normality implicit in the analysis of variance. *Journal of the American Statistical Association*. 1937;32(200):675–701. https://doi.org/10.1080/01621459.1937.10503522

[29] Wilcoxon F. Individual comparisons by ranking methods. *Biometrics Bulletin*. 1945;1(6):80–83. https://doi.org/10.2307/3001968

[30] Holm S. A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*. 1979;6(2):65–70.