Student Trajectories in Higher Education

Curriculum Structure, Causal Inference, and Simulation for Policy Design

Hugo Roger Paz

Student Trajectories in Higher Education

Curriculum Structure, Causal Inference, and Simulation for Policy Design

Student Trajectories in Higher Education

Curriculum Structure, Causal Inference, and Simulation for Policy Design

Hugo Roger Paz, PhD
Professor and Researcher
Faculty of Exact Sciences and Technology
National University of Tucumán
San Miguel de Tucumán, Argentina

First edition, 2026.

Independently published.

ORCID: 0000-0003-1237-7983

Editorial transparency note

The author used linguistic and editorial assistance tools during manuscript preparation. The conceptual framework, methodological design, interpretation, and final responsibility for the work remain entirely the author's own.

To my family

Preface

This book addresses student trajectories in higher education as a structural and temporal problem rather than as a purely individual one. Its particular focus is on strongly sequenced programmes, where progression is shaped not only by student characteristics and academic performance, but also by curriculum topology, institutional regulation, and contextual disturbance.

The volume develops CAPIRE as an integrated analytical framework for the study of curriculum-constrained trajectories. The framework combines leakage-aware data design, curriculum graph modelling, heterogeneity analysis, causal estimation, and agent-based simulation within a single decision-support architecture. The aim is not simply to improve risk detection, but to make structural mechanisms visible and institutionally actionable.

The empirical illustrations are drawn from a public Civil Engineering programme in Argentina. They are used as an operational case through which the framework is specified, tested, and discussed. The intention is not to claim direct empirical generalisability from one institutional setting to higher education as a whole, but to offer a transferable analytical logic that can be adapted to different curricula, data environments, and regulatory regimes.

The chapters are organised progressively. The opening chapters establish the conceptual and methodological problem. The middle chapters develop the technical components of the framework. The later chapters address integration, institutional use, governance, and transferability.

This book is intended for researchers, institutional analysts, and academic decision-makers working in higher education, engineering education, learning analytics, curriculum analysis, and institutional policy. Its central claim is that meaningful intervention in student trajectories requires moving beyond flat administrative indicators and prediction-only approaches towards a structurally informed, temporally valid, and mechanism-oriented analytical model.

Acknowledgements

This book was developed through a long process of research, writing, and methodological refinement carried out in parallel with teaching and institutional responsibilities. I am grateful to the colleagues, students, and academic environments that, over many years, helped make these questions both intellectually urgent and empirically concrete.

I also acknowledge the broader body of scholarship in higher education, learning analytics, causal inference, network science, and complex systems that made this work possible. The framework developed here builds on those traditions while attempting to connect them in a way that is analytically coherent and institutionally useful.

Special thanks are due to the institutions and academic settings that sustained the empirical and conceptual work underlying this volume, and to all those who, directly or indirectly, contributed to the long maturation of its central ideas.

Any errors, limitations, or interpretive choices that remain are entirely my own.

Contents

Chapter 1
Introduction: The Challenge of Student Trajectories in Engineering Education

ABSTRACT

Student attrition in engineering programmes is a persistent and poorly understood phenomenon. Decades of retention research and learning analytics have focused on individual-level explanations - motivation, preparation, financial stress, social integration - while treating the curriculum and its regulatory framework as background context. Where curriculum structure has been analysed formally, it has been analysed at the level of the programme rather than the individual student. This chapter argues that such framing is insufficient for curriculum-constrained programmes, where tightly structured prerequisite architectures, high-stakes gateway courses, and strict regularity rules actively shape what is possible for students at every stage of their trajectory. In such programmes, the organisation of the curriculum itself is one of the principal mechanisms through which trajectories are produced, amplified, and foreclosed. The chapter opens by situating engineering attrition as a structural rather than an incidental problem, drawing on the case of a public engineering programme operating under open admission, dense prerequisite dependencies and pronounced economic volatility. It then traces the limits of existing student-level explanations and introduces the shift from individual-centred analysis to structural diagnosis. The chapter concludes by previewing the book's analytical architecture - the CAPIRE framework - and the five interconnected components that together make curriculum-constrained trajectories visible, explainable, and policy-relevant. The central claim is that students do not navigate a neutral academic space; they move through a designed environment whose topology determines what is possible and what is blocked, and shapes how difficult recovery becomes as time passes.

1.1 Why Student Trajectories Matter in Engineering Education

Every year, engineering programmes with long, densely sequenced prerequisite structures lose a substantial proportion of their enrolled students. Some leave before completing their first semester. Others persist for years - passing some courses, failing others, accumulating delays - and then disappear quietly from institutional records without any discernible departure event. Still others graduate, but only after trajectories so extended and interrupted that the programme's nominal duration has long ceased to serve as a realistic benchmark. This is not a fringe phenomenon. In many public higher-education systems, completion is far from assured, and time to degree diverges substantially from nominal programme length; participation and completion remain sharply unequal across regions (Tight, 2020; UNESCO, 2022). In Argentine public universities, extended time to degree and low on-time completion are long-standing and well-documented patterns (Giovagnoli, 2002; García de Fanelli, 2014). In the curriculum-constrained engineering programme examined in this book, completion within the nominal timeframe is the exception rather than the rule (Paz, 2023a, 2023b).

The familiar response is to study the students. Decades of retention research have examined motivation, prior academic preparation, social integration, financial stress, and institutional engagement. Learning analytics has extended this tradition by adding predictive power: with sufficient data and appropriate methods, models can identify students most likely to withdraw before they actually do (Baker & Inventado, 2014; Romero & Ventura, 2020). These contributions are genuine and consequential. Institutions that have invested in early-warning systems have in many cases improved their capacity to identify vulnerable students and direct support resources more efficiently.

Yet these approaches share a premise that this book questions: that the primary object of analysis is the individual student, and that the curriculum, the institutional regulatory framework, and the structural environment are, at most, contextual background. This book proceeds from a different starting point. It argues that in engineering programmes with tightly structured prerequisite architectures - long dependency chains, high-stakes gateway courses, strict regularity rules, and limited tolerance for deviation from the nominal timeline - the organisation of the curriculum itself is one of the principal mechanisms through which student trajectories are shaped. Students do not navigate a neutral academic space. They move through a designed environment whose topology determines what is possible and what is blocked, and shapes how difficult recovery becomes as time passes.

Box 1.1 What is a student trajectory?
In this book, a student trajectory refers to the sequence of academic events - enrolments, course outcomes, delays, interruptions, and progressions - through which a student moves from entry to a final outcome (graduation, departure, or prolonged enrolment). Trajectories are not merely chronological records; in curriculum-constrained programmes, they are structured processes shaped by prerequisite dependencies, regulatory rules, and contextual conditions. The same sequence of events may have very different consequences depending on when it occurs and which courses it involves. Analysing trajectories therefore requires attending to timing, position within the curriculum network, and the institutional rules that govern what is possible at each point.
Note: The term is used descriptively, not normatively. A long trajectory is not per se a failed one; the concern is with structural fragility - the degree to which a trajectory has become difficult to recover from, regardless of its current administrative status.

The argument is not that individual characteristics are irrelevant. Prior preparation, economic resources, work obligations, and social integration all matter, and the evidence for their influence is well established. The point is that these factors operate within a structured environment, and that the structure itself - its prerequisite graph, its bottleneck courses, its regulatory thresholds - amplifies some difficulties and absorbs others in ways that cannot be understood by studying students alone. An analytical framework adequate to this problem must treat the curriculum as an object of analysis, not as background.

1.2 Beyond Student-Level Explanations

The foundational literature on student departure established several insights that remain indispensable. Spady (1970) and Tinto (1975, 1993) demonstrated that departure is not simply a matter of individual failure or insufficient resolve; it is shaped by the quality of academic and social integration - by whether students experience a genuine sense of belonging and institutional recognition. Bean (1980) and Bean and Metzner (1985) extended the picture by showing that, particularly for non-traditional students, external pressures - work obligations, family responsibilities, financial instability, commuting burdens - can be more consequential than campus-based integration. Astin (1993) argued that active involvement with the institutional environment is itself a driver of persistence, not merely a symptom of prior preparation. Pascarella and Terenzini

(2005) synthesised three decades of evidence to show that how higher education affects students depends on the cumulative interaction between individual characteristics and the institutional experience as a whole.

These contributions remain the intellectual foundation of the field, and this book builds on them. They share, however, a structural blind spot that becomes increasingly visible in curriculum-constrained settings: the curriculum itself receives remarkably little analytical attention as an independent source of trajectory outcomes. In the dominant retention models, the curriculum appears as a backdrop - the environment in which academic integration happens, the occasions on which grades are earned - rather than as an organised structure with topological properties that may systematically advantage or disadvantage students independently of their preparation, motivation, or institutional belonging.

The rise of learning analytics partially addressed this gap, but primarily by improving the detection of at-risk students rather than by clarifying what produces risk in the first place. Baker and Inventado (2014), Leitner et al. (2017), Romero and Ventura (2020), and a substantial literature on educational data mining demonstrated that administrative records and institutional databases can be combined to predict dropout risk with considerable accuracy. Early-warning systems became a central institutional use case: identify students showing patterns associated with eventual departure and intervene before damage becomes irreversible (Macfadyen & Dawson, 2010; Delen, 2010). These applications have been useful. They have moved institutions away from the reactive practice of investigating attrition only after it has already occurred.

Predictive accuracy, however, is not the same as explanatory power. A model can learn to identify at-risk students by detecting patterns associated with departure without clarifying what produces those patterns. A blocked-credit indicator may predict withdrawal without telling an institution whether the underlying problem lies in student preparation, course difficulty, prerequisite architecture, assessment design, regulatory rules, or some interaction among these. The most consequential limitation is analytical: predictive models trained on historical data tend to reproduce the institutional logic embedded in that data. If a curriculum has functioned as a mechanism of stratification - consistently concentrating failure among students with less prior preparation or less economic resilience - a model trained on that history may achieve accuracy precisely by learning to reproduce the distribution of disadvantage

(O'Neil, 2016; Slade & Prinsloo, 2013). Prediction, in that case, describes the problem without illuminating it.

The argument developed in this book is not that prediction is without value. It is that prediction placed at the centre of institutional analytics leaves the most important questions unanswered: which mechanisms produce vulnerable trajectories; where in the programme structural risk concentrates; whether the administrative rules that claim to uphold academic standards are also generating avoidable exclusion; and what would change if specific curricular features or regulatory thresholds were redesigned. These are questions of mechanism and structure, and they require causal reasoning, formal structural representation, and simulation-based exploration - not only pattern detection (Hernán & Robins, 2020; Pearl, 2009).

1.3 An Illustrative Case of Structural Fragility

To make the argument concrete before making it theoretical, consider a student entering an engineering programme with adequate secondary school preparation, no particular socioeconomic disadvantage, and a genuine intention to graduate. In the first year, two mathematically intensive foundational courses run simultaneously. One is passed on the first attempt; the other is not. The failed course, it turns out, is a prerequisite for four subsequent subjects distributed across the following three semesters. The student cannot enrol in those four courses immediately. The available options for the following term contract sharply: fewer courses, fewer credits, less forward momentum.

A temporary disruption then compounds the situation. It might be financial - the student takes a part-time position to help cover household expenses during a period of rising prices. It might be institutional: a prolonged dispute between academic staff and the university compresses the examination calendar, collapsing two assessment periods into one. Or it might be personal, a family obligation that reduces available study time for six weeks at a structurally sensitive moment. In a programme with distributed prerequisites, multiple paths through the curriculum, and generous provisions for re-entry after setbacks, this combination of events would be inconvenient but manageable. In a programme with dense prerequisite dependencies and strict regularity requirements, it is not. One unmet prerequisite feeds into another. The student's structural distance to graduation - not the number of outstanding credits, but the length and rigidity of the path that must still be traversed - begins to extend in ways that are not visible in any standard administrative indicator.

What is analytically significant here is what the institutional record cannot see. Standard institutional records - grades, credit counts, regularity status - will register this student's growing fragility only poorly, and only late. The blocked credits appear in administrative systems, but the structural meaning of those blocks - which pathways they close, how deep the downstream dependency chain extends, how far the graduation horizon has shifted - remains invisible without a formal representation of the curriculum as a network. By the time the student appears in a risk score or an adviser's caseload, the trajectory may already be very difficult to reverse.

Multiply this across an entire cohort - with heterogeneous preparation levels, varying degrees of economic exposure, and differential access to support - and the aggregate outcome begins to look less like a collection of individual failures and more like a predictable consequence of a particular structural design. This is the distinction the analytical framework developed in this book is built to make legible.

1.4 From Prediction to Structural Explanation

The shift proposed in this book is not a rejection of predictive analytics but a repositioning of its role within a broader analytical architecture. Prediction remains operationally useful: institutions need to identify patterns of risk, allocate support resources, and detect emerging vulnerabilities before they become irreversible. But in curriculum-constrained programmes, the deeper challenge is to understand how trajectories are shaped over time by the interaction of individual heterogeneity, curricular architecture, institutional regulation, and contextual disturbance. Achieving that understanding requires capacities that predictive models, taken alone, do not provide.

The first is temporal integrity. Many educational datasets are assembled retrospectively, and many models are trained without disciplined attention to what information was actually available at each decision point. When a model uses features that post-date the observation window it is intended to inform - a form of data leakage that is, in the assessment developed in Chapter 5, considerably more common in applied educational work than the published accuracy figures suggest - it will appear more accurate than it is when deployed in practice (Kaufman et al., 2012). Leakage is not a minor technical nuisance; it undermines the validity of any analytics system intended to support prospective institutional action. A framework serious about decision support must design the data layer, from the outset, to prevent it.

The second is structural visibility. A student's academic status as recorded in administrative systems - courses passed, courses failed, credits accumulated - conveys outcomes but almost nothing about position. Position, in the sense relevant here, means something specific: which courses are currently accessible, which are blocked, how long the dependency chain to graduation is, and how much structural slack remains for absorbing further setbacks. None of that is recoverable from a flat administrative record. It requires representing the curriculum as a network - a directed graph in which nodes are courses and edges are prerequisite dependencies (Heileman et al., 2018; Molontay et al., 2020) - and then deriving from that representation structural indicators defined at the level of the individual student at a given point in time (Paz, 2025b).

The third is causal reasoning. Predicting which students will leave is a different problem from understanding why trajectories become fragile and which institutional changes would alter their distribution. The second question requires the language of mechanism: what do specific curricular features, regulatory rules, or macro-level shocks do to trajectory outcomes, for which kinds of students, and under what conditions? Answering it requires moving from association to mechanism - a transition that involves explicit causal assumptions, appropriate estimation methods, and interpretive discipline (Hernán & Robins, 2020; Pearl, 2009; Chernozhukov et al., 2018).

Once these three requirements are in view, a fourth becomes apparent: the capacity to explore what might happen under conditions that have not yet occurred. Institutions contemplating regulatory reform, curricular redesign, or changes in support provision need some way to reason about consequences before implementing changes on live cohorts. Agent-based modelling provides that capacity - not by predicting the future with precision, but by making explicit the assumptions about how students, curriculum structures, and institutional rules interact, and by examining what those interactions produce over time under alternative policy configurations (Epstein & Axtell, 1996; Miller & Page, 2007).

Temporal integrity, structural visibility, causal reasoning, and simulation-based foresight are the four analytical pillars on which the framework developed in this book is built.

1.5 What CAPIRE Adds

CAPIRE - the acronym expands to Curriculum-Aware Pipeline for Institutional Research and Education - is the organising framework of this book. Its purpose is not to displace existing retention scholarship but to connect strands of analysis that are too often pursued in isolation: temporally valid data design, explicit representation of curriculum structure, analysis of trajectory heterogeneity, causal estimation, and simulation-based policy testing. The integration is not cosmetic. Each component addresses a genuine limitation in how student trajectories are typically studied, and the components constrain and enrich one another in ways that matter analytically.

The data layer establishes the temporal foundation. It implements a leakage-aware approach to constructing analytical records - a disciplined procedure for defining observation windows, separating what was available at each decision point from what was revealed only later, and organising variables into multilevel layers that distinguish pre-entry background from trajectory dynamics and contextual exposure. Without this layer, later analytical sophistication rests on an unstable empirical basis.

The curriculum graph translates the programme's prerequisite structure into a formally analysable representation. Once the curriculum is encoded as a directed acyclic graph, it becomes possible to compute student-level structural indicators: blocked credits (how many credits the student currently cannot access due to unmet prerequisites), backbone completion (progress through the structurally most critical courses), bottleneck exposure (position relative to courses that mediate the most downstream access), and graph-based distance to graduation (the length and rigidity of the remaining dependency path). These indicators make analytically visible what administrative records leave implicit.

The heterogeneity layer recognises that student populations are not homogeneous and that analytical frameworks calibrated to an average student will produce interventions that fit few students well. Using dimensionality reduction and density-based clustering applied to the leakage-aware feature space, CAPIRE identifies recurring trajectory archetypes - not as fixed psychological types, but as structural patterns of progression that recur across cohorts and carry different implications for risk, responsiveness to support, and sensitivity to regulatory change.

The causal layer addresses questions of mechanism. Once structural position and trajectory heterogeneity are visible, the natural next question is which factors move outcomes rather than merely correlate with them. Double machine learning estimation allows this question to be asked with explicit assumptions, flexible confounding control, and interpretive discipline - producing mechanism-oriented findings rather than correlation rankings (Chernozhukov et al., 2018).

The simulation layer closes the pipeline. Agent-based modelling allows the analyst to examine what would happen under policy alternatives that have not yet been implemented - regulatory changes, curricular redesign, targeted support for specific trajectory archetypes, compound scenarios. Not as prediction, but as structured hypothesis: a disciplined way of reasoning about interaction effects and unintended consequences before they affect real students.

The distinctiveness of CAPIRE lies less in any single component than in their integration. The data layer informs the graph; the graph informs heterogeneity analysis; heterogeneity informs causal estimation; causal estimation informs simulation. The whole is not simply the sum of its parts. It is a pipeline in which each layer is constrained and enriched by the others. Figure 1.1 summarises this integrated pipeline and shows how each analytical layer informs the next, culminating in institutionally usable decision outputs.

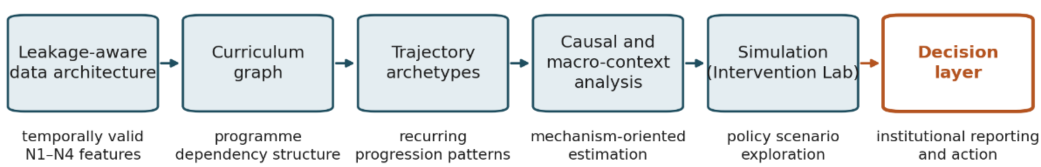

Figure 1.1 The CAPIRE analytical pipeline. The five components - leakage-aware data architecture, curriculum graph, trajectory heterogeneity, causal estimation, and agent-based simulation - function as an integrated sequence in which each layer informs and disciplines the next. Outputs feed into the institutional decision layer (Chapter 9).

1.6 Scope, Transferability, and Responsible Use

The empirical illustrations in this book are drawn from a single institutional setting: a public Civil Engineering programme at the National University of Tucumán, Argentina. That context was not selected because it is representative of engineering education globally, nor because its conditions are unusual enough to be merely exotic. It was selected because its features - open-access admissions without entry examinations, long and densely structured prerequisite chains, strict regularity rules, pronounced economic instability, and recurrent academic staff industrial action - make the structural mechanisms under analysis unusually visible. In more forgiving institutional environments, the same mechanisms should be expected to operate with smaller and less easily detected effects. In this one, they are difficult to miss.

The framework is therefore grounded in a specific context but is not confined to it. The analytical logic of CAPIRE - temporal integrity, structural representation, heterogeneity analysis, causal reasoning, simulation - travels across settings because it responds to properties of the problem rather than properties of one institution. What does not travel, and should not be expected to, are the empirical specifics: the particular bottleneck courses, the specific regulatory thresholds, the magnitude of macro-shock effects, the shape of the archetype distribution. These require local calibration in each new setting.

Box 1.2 What does 'illustrative case' mean in this book?

Throughout this book, the Civil Engineering programme at the National University of Tucumán, Argentina, functions as an illustrative case. This framing has a precise methodological meaning: the case is used to show that the analytical architecture of CAPIRE can be operationalised coherently on real institutional data, and to demonstrate concretely how structural mechanisms operate. It is not used to establish universal empirical regularities or to claim that findings from one programme generalise directly to all engineering education systems.

Readers working in other institutional contexts will need to calibrate the framework's components - graph metrics, archetype definitions, causal models, simulation parameters - to their own curricula, data infrastructures, and regulatory environments. The case illustrates what is possible; it does not prescribe what will be found. Chapter 13 addresses transferability in detail.

This distinction between transferable analytical logic and locally calibrated parameters is a persistent methodological commitment of this book. It is what allows the framework to claim broader relevance without overreaching. Institutions in other national and regulatory contexts that face curriculum-constrained attrition problems are not being invited to replicate the Tucumán case. They are being offered an analytical architecture that can be configured to their own curricula, their own data environments, and their own institutional rules. Portability here does not mean context-independence. It means that contextual adaptation is a visible and disciplined part of the analytical process rather than an afterthought.

1.7 Governance, Transparency, and Responsible Institutional Analytics

A framework intended to inform consequential institutional decisions about student trajectories cannot treat the politics of evidence as an afterthought. Once analytics begins to influence which students receive support, which programmes are restructured, and how vulnerability is classified, it enters the domain of institutional power. The question of how that power is exercised - transparently or opaquely, with accountability or without - is not separable from the technical quality of the analysis.

Three principles run through the entire framework. Transparency requires that institutional users be able to trace how any analytical output was produced: what data were used, what observation window was applied, what modelling assumptions were made, and where the limits of the result lie. Interpretability requires that outputs be explainable in ordinary institutional language, not only in statistical terms - because an output that cannot be explained cannot be contested, and an output that cannot be contested cannot be corrected (Adadi & Berrada, 2018; Lundberg & Lee, 2017). Accountable human judgement, finally, means that analytics support decision-making rather than replace it: the framework is built as decision infrastructure, not as automated authority.

These considerations are not peripheral to the framework. They belong to its design logic. Chapter 12 develops them fully in relation to governance, transparency, and responsible institutional analytics. They are introduced here because governance is not a destination reached at the end of a methodological journey. It is a condition that shapes every analytical decision made along the way.

1.8 How This Book Is Organised

The book develops its argument progressively, with each chapter building on and disciplining the preceding one. Chapter 2 examines the dominant approaches to student retention and learning analytics, identifying the specific analytical gaps - structural blindness, temporal invalidity, and the conflation of prediction with explanation - that motivate the CAPIRE architecture. Chapter 3 establishes the conceptual foundations: what it means to treat a student trajectory as a curriculum-constrained process, and how complexity theory provides analytical vocabulary without becoming an explanatory straitjacket. Chapter 4 introduces CAPIRE as an integrated architecture and explains how its components relate to one another.

Chapters 5 through 8 develop the methodological core. Chapter 5 addresses data architecture and leakage-aware design - the epistemic foundation on which everything else rests. Chapter 6 presents curriculum structure as a network, developing the graph-based indicators that make structural position analytically visible. Chapter 7 develops the analytical layer, combining descriptive, archetype-based, and causal approaches. Chapter 8 introduces the simulation layer, showing how agent-based modelling converts empirical findings into structured policy foresight. Chapter 9 integrates these components into a coherent decision-support pipeline.

Chapter 10 presents the institutional case study - not as a proof of generality but as a demonstration that the framework functions as a unified architecture. Chapter 11 turns from analysis to policy, examining how evidence can inform intervention design without overclaiming certainty. Chapter 12 addresses governance, transparency, and the responsibilities that accompany consequential institutional analytics. Chapter 13 discusses transferability, contextual calibration, and the broader implications of the framework for engineering education systems. Chapter 14 concludes by drawing together the book's main contributions, its acknowledged limitations, and the directions it opens for future work. Table 1.1 provides a compact overview of the main analytical traditions discussed in the next chapter, their primary focus, and the way CAPIRE relates to each of them. The volume also includes technical appendices that document selected design, implementation, and reproducibility elements supporting the main argument.

Table 1.1 Prevailing analytical approaches to student retention, their primary focus, treatment of curriculum structure, and relationship to the CAPIRE framework

Approach	Primary analytical focus	Treatment of curriculum structure	Relationship to CAPIRE
Classical retention models (Tinto, 1975; Bean, 1980; Pascarella & Terenzini, 2005)	Social and academic integration; individual persistence decisions	Implicit background - the site where integration occurs, not an independent analytical object	Foundational insights on the relational and processual nature of departure; CAPIRE adds structural representation
Learning analytics and early-warning systems (Baker & Inventado, 2014; Romero & Ventura, 2020)	Statistical patterns of risk; classification of at-risk students	Absent or encoded as flat credit counts and subject labels; dependency topology not modelled	Provides triage utility; CAPIRE reframes prediction as one layer within a causal and structural pipeline
Educational data mining and machine learning (Delen, 2010; Leitner et al., 2017)	Predictive accuracy; feature importance; model performance metrics	Curriculum features appear as variables, not as a network; temporal sequencing typically absent	Methodological discipline on validation; CAPIRE adds leakage-aware temporal design and structural indicators
Causal inference approaches (Hernán & Robins, 2020; Pearl, 2009)	Mechanism identification; treatment effect estimation under explicit assumptions	Structural covariates may be included but curriculum topology is rarely formalised as a graph	Causal layer within CAPIRE; combines double machine learning with graph-derived structural features
CAPIRE (this volume)	Curriculum-constrained trajectories as emergent processes; structural, temporal, causal, and contextual analysis	Formally represented as a directed acyclic graph; student-level structural indicators derived from topology	Integrates preceding traditions within a coherent pipeline oriented towards institutional decision support

Note. This table summarises the principal analytical traditions addressed in Chapter 2. The characterisation of each approach reflects its dominant orientation in the applied literature; individual studies within each tradition may deviate from the patterns described.

CONTINUE READING

STUDENT TRAJECTORIES IN HIGHER EDUCATION

Choose the store you prefer. Click or tap a panel to open it.

BUY THE EBOOK ON AMAZON

Kindle edition. Available instantly.

BUY THE PAPERBACK ON AMAZON

Print edition. Available worldwide.

VIEW ALL STORES ON BOOKS2READ

Available stores and formats in one public link.

Student trajectories do not unfold in a neutral academic environment. They are shaped by curriculum structures, prerequisite dependencies, institutional regulations, temporal accumulation, and contextual disturbance.

Student Trajectories in Higher Education presents CAPIRE, a curriculum-aware analytical framework that combines leakage-aware data architecture, curriculum graph modelling, trajectory heterogeneity, causal inference, and agent-based simulation.

Rather than asking only which students are at risk, the framework asks how institutional structures produce, amplify, or constrain different trajectories - and how analytical evidence can support better-informed policy decisions.

Drawing on a longitudinal Civil Engineering case from Argentina, the book demonstrates how this architecture can be implemented on real institutional data while maintaining a clear distinction between context-specific empirical findings and transferable analytical principles.

Designed for researchers, institutional analysts, and academic decision-makers, the book offers a rigorous framework for moving from prediction towards structural understanding, causal reasoning, and responsible policy simulation.