

Presentation Companion

A Concise Chapter Digest — Policy by the Numbers

A medium-density summary of all 12 chapters — more detail than a one-line summary, far less than the full chapters. Built for briefings, board packets, and presentation prep.

Chapter 1: The Fundamentals of Educational Policy Modeling

Defines the role of modeling in modern education, shifting from intuition-based policy to a predictive, Logistic Regression framework. The Mathematical Model: $\ln(P/(1-P)) = \beta_0 + \beta_1R + \beta_2I + \beta_3A + \beta_4S + \epsilon$.

Chapter 2: Why the Logit Function? Choosing the Right Tool for the Job

Answers the question every technical skeptic eventually asks — why logistic regression, not a fancier machine learning model? The logit function keeps predictions inside a sensible 0–100% range, and its coefficients translate into odds ratios a board, a parent, and a reporter can all repeat back accurately.

Chapter 3: Finding and Preparing Your Data

A roadmap for locating, verifying, and preparing public data. Highlights: normalization (Z-scores), subgroup representation, trend analysis. Sourcing quality data requires a full lifecycle of acquisition, cleaning, and preparation.

Chapter 4: Designing and Implementing Your First Model

Scenario Case Study (85% Graduation Target): $R=0.70$, $A=0.92$, $I=0.85$, $S=0.40$; $z \approx 1.475$; $P \approx 0.814$. Corrected sensitivity finding: attendance alone reaches only 82.5%; reaching 85% requires a combined attendance-plus-academic-support push.

Chapter 5: Scenario Planning: Adapting Models for Diverse Districts

The Scenario Matrix maps district position to a named strategic quadrant. Equity Gap Ratio = $P(\text{subgroup}) \div P(\text{total})$; below 0.8 signals floor resistance.

Chapter 6: Communication Strategies: Translating Data to Policy

Audience Segmentation: school boards (ROI), parents (student impact), community (transparency). The Green Zone Narrative simplifies results into color-coded thresholds.

Chapter 7: Policy Guide — From Prediction to Implementation

Case Study — The Equity Lift: baseline $P \approx 58\%$ (Tier 3, $k_s = 1.0$) to $P \approx 73\%$ (Tier 2, $k_s = 1.20$). Fiscal Efficiency Mapping ties k_s ranges to ROI tiers for grant justification.

Chapter 8: From One District to Every District

Running the identical model across districts of different sizes produces very different P values. Statewide P is an enrollment-weighted average, never a simple unweighted one.

Chapter 9: Formal Specification of the Success Probability Model

Marginal Effects: $\partial P / \partial x = \beta \cdot P(1-P)$, not β alone. Hierarchical Model (GLMM) formalizes the state-district structure via empirical-Bayes partial pooling.

Chapter 10: Year Two — Did the Model Actually Work?

Predicted vs. actual: general cohort 85% predicted, 83.1% delivered; EL cohort 73% predicted, 68.4% delivered. Confusion matrix reveals specificity (55%) weaker than raw accuracy (83%) suggests.

Chapter 11: Governance, Accountability, and the Limits of the Model

Four required roles: Model Owner, Data Steward, Equity Reviewer, Human-in-the-Loop Lead. Published, non-negotiable limit on what the model may decide alone.

Chapter 12: From State Dashboard to Federal Accountability

Maps R, A, I, S, V directly onto ESSA's five required accountability indicators. NAEP is framed as an external calibration check, not a data source.

THE ONE-SLIDE VERSION *The problem: end-of-year data arrives too late. The tool: one formula, five variables, and a deliberate choice of logistic regression over a fancier model. The equity mechanism: k_s quantifies the lift from targeted support. The proof: it scales from one district to a whole state to a rigorous statistical foundation, and it's been checked against real outcomes. The guardrails: named governance roles and a published list of decisions the model may never make alone. The one sentence: diagnostic, not deterministic.*