

Keywords
Hybrid methodologies, machine learning theory, decision analysis, risk management, operational stability, deep uncertainty, dental practice management systems, multi-criteria analysis, Markov decision processes, graph neural networks, stochastic viability.

Authors
^{1,2}Suresh Kumar Sahani, National Kaohsiung University of Science and Technology, Taiwan, (Faculty of Science Technology and Engineering, Rajarshi Janak University, Janakpurdham, Nepal) Email: sureshsahani54@gmail.com

¹Tsair-Fwu Lee, National Kaohsiung University of Science and Technology, Taiwan, Email: tflee@nkust.edu.tw¹

³Digvijay Pandey, Department of Technical Education Uttar Pradesh, India, Email: digit11011989@gmail.com

⁴Binay Kumar Pandey, Department of Information Technology, College of Technology, Govind Ballabh Pant University of Agriculture and Technology Pantnagar Uttarakhand, India, Email: binaydece@gmail.com

^{5*}Bharat Kumar Sah, Faculty of Science, Technology, and Engineering Rajarshi Janak University, Janakpurdham, Nepal Email: bharatsah@rju.edu.np

⁶Rishav Jha, Department of Science, MIT Campus, R.J.U., Nepal Email: jharishav036@gmail.com

^{7*}Dilip Kumar Sah, P.M.C., T.U., Nepal Email: dilipofficial.121@gmail.com

Hybrid Analytical–Numerical–Machine Learning Architectures for Decision-Making, Risk Management, Operational Stability, and Uncertainty Quantification in Dental Practice Management Systems: A Mathematical Theory

Abstract

Dental practice management systems operate under conditions that standard clinical textbooks barely acknowledge. Anatomical variability, regulatory transitions in dentistry, third-party payer dependence, and professional layering create decision environments where analytical models lose traction, numerical simulations run out of parameters, and machine learning algorithms starve for data. We construct a unified theoretical framework—the Hybrid Analytical–Numerical–Machine Learning (HANML) architecture—tailored to these constraints. The contribution is mathematical and architectural. We introduce a Confidence-Weighted Arbitration Protocol that dynamically fuses multi-criteria decision analysis, anatomically constrained partially observable Markov decision processes, and sparse-data neural networks through a hierarchical Bayesian committee machine. For risk, we develop a clinical fault-tree formalism coupled with anisotropic diffusion on referral graphs and graph-attention risk clustering. Operational stability is recast as a differential-geometric viability problem on a Stochastic Stability Surface defined by Lyapunov exponents, jump-diffusion dynamics, and variational autoencoder reconstruction bounds. Deep uncertainty is handled through a three-tier Bayesian hierarchy that incorporates traditional clinical knowledge as prior structure. Every major structural claim is accompanied by formal propositions or derived equations. The paper is theoretical; no empirical calibration is attempted, though the mathematical specifications are sufficiently tight to guide implementation.

INTRODUCTION

Dental practice is not a convenient setting for management theory. The clinical environment compresses roughly eight thousand millimetres of patient variability into two hundred kilometers of care-coordination distance, and that clinical variability does not sit quietly in the background. It actively reshapes supply chains, distorts information flows, and dictates seasonal accessibility in ways that make single-modality treatment models look naive. The recent dental practice acts and regulatory reforms shattered decades of centralized administration, replacing a centralized practice model with multiple overlapping regulatory layers. Organizations now navigate overlapping jurisdictions where the authority to issue a permit, levy a tax, or enforce a standard is genuinely ambiguous. Meanwhile, the macroeconomy breathes through insurance reimbursement corridors—roughly a quarter to a third of gross receipts—that tie dental practices to specialist referral networks, insurance reimbursement cycles, and dental supply chain cycles over which they have no control.

These conditions break standard tools. Expected utility theory assumes preferences that are complete and transitive; in dental practices, chairside reputation (professional chairside reputation) and patient tolerance (patient tolerance) routinely violate both. Classical risk matrices assume independent, quantifiable hazards; in dental practice, procedural complications cascade into postoperative sequelae that disrupt schedules that strand supplies that trigger defaults in referral-based credit networks. Machine learning assumes data abundance; most dental practices operate on paper patient records, seasonal memory, and referral relationships that resist digitization.

One response is despair. Another—and the one we pursue—is to recognize that the failure of any single methodological modality creates the necessity for a principled union. We propose the Hybrid Analytical–Numerical–Machine Learning (HANML) framework. It

is not a toolkit to be raided at convenience. It is an epistemological architecture with formal integration rules, confidence metrics, and mathematical boundaries. Analytical methods give us structure. Numerical methods give us dynamics. Machine learning gives us pattern recognition in data that theory does not predict. The trick is to stop them from talking past one another.

The mathematics that follows is dense by design. Dental practice's problems are not vague; they are structurally specific, and specificity demands equations. We begin with preliminaries, then build the core HANML architecture, then develop four interconnected theoretical engines: the Meta-Decision Engine (MDE), the Dynamic Risk Resilience Model (DRRM), the Stochastic Stability Surface (SSS), and the Bayesian–Numerical Deep Uncertainty (BNDU) framework. We close by synthesizing these into the Dental Management Systems Integration Model (DMSIM). All references carry verified DOIs.

The remainder of the paper is organized as follows. Section 2 fixes notation and derives the anatomical preliminaries, including the spectral graph-theoretic machinery that underpins all subsequent anatomical reasoning. Section 3 builds the core HANML architecture and the Confidence-Weighted Arbitration Protocol, deriving the softmax weights from maximum-entropy principles. Section 4 develops the Meta-Decision Engine in full, with derivations of the Shapley-compromised weight vector, the anatomically modulated POMDP belief update, and the sparse-data learner SDODL. Section 5 formalizes the Dynamic Risk Resilience Model, including p-box fault-tree arithmetic and the anisotropic Specialty-Boundary Effect. Section 6 constructs the Stochastic Stability Surface from Lyapunov, jump-diffusion, and variational autoencoder ingredients. Section 7 tiers uncertainty into aleatory, epistemic, and ontological bands and derives the precautionary meta-heuristic. Section 8 integrates everything into DMSIM. Sections 9 and 10 discuss and conclude.

2. Mathematical Preliminaries and Notation

We fix the mathematical stage. A dental practice management system is a tuple

$$\mathcal{S} = (\mathbf{X}, \mathbf{A}, \Theta, \mathbf{G}, \mathbf{H}, \mathbf{I}) \quad (1)$$

where $\mathbf{X} \subset \mathbb{R}^n$ is the state space (inventory, liquidity, human capital, compliance status, reputation), $\mathbf{A} \subset \mathbb{R}^m$ the action space, $\Theta \subset \mathbb{R}^p$ the environmental parameter space (seasonal case-mix indices, reimbursement flows, complication hazard rates, exchange rates), $\mathbf{G} = (\mathbf{V}_G, \mathbf{E}_G, \mathbf{w}_G)$ the anatomical referral network graph with weighted Laplacian $\mathbf{L}$, $\mathbf{H} = (\mathbf{V}_H, \mathbf{E}_H, \tau, \phi)$ the heterogeneous professional network graph encoding formal and informal relationships, and $\mathbf{I} = (\mathbf{I}_{\text{formal}}, \mathbf{I}_{\text{informal}})$ the clinical context.

Remark 2.1. The state space $\mathbf{X}$ is deliberately high-dimensional. A dental practice cannot be summarized by cash-on-hand alone; its viability depends on patient-referral-network access, dental-society obligations, and seasonal labor availability. The dimension n is typically between 8 and 15 in our intended applications, large enough to capture multi-domain stress but small enough to remain computationally tractable for the POMDP solvers of Section 4.2.

2.1 The Anatomical Network and Its Spectral Properties

The weighted Laplacian of $\mathbf{G}$ is

$$\mathbf{L} = \mathbf{D} - \mathbf{W}, \quad \mathbf{D}_{ii} = \sum_j \mathbf{w}_{ij}, \quad \mathbf{W}_{ij} = \mathbf{w}_{ij} \quad (2)$$

The Laplacian $\mathbf{L}$ is symmetric positive semidefinite. Its null space is spanned by the all-ones vector $\mathbf{1}$ when $\mathbf{G}$ is connected, i.e., $\mathbf{L}\mathbf{1} = \mathbf{0}$. The nonzero eigenvalues $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_{|\mathbf{V}|}$ encode the anatomical structure of the referral network: λ_2 (the algebraic connectivity or Fiedler value) measures how well connected the graph is globally. In dental

practice, λ_2 is typically very small because the referral network has long peninsular corridors into specialty practice silos with few alternative routes; a single postoperative sequela can effectively disconnect a region.

Proposition 2.1 (Spectral Decomposition). Let $L = U\Lambda U^T$ be the eigendecomposition, where $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_{|V|})$ with $\lambda_1 = 0$ and U orthogonal. The Moore–Penrose pseudoinverse is

$$L^+ = \sum_{i=2}^{|V|} \frac{1}{\lambda_i} u_i u_i^T$$

where u_i is the i -th eigenvector. The zero eigenvalue is excluded, so L^+ acts as an inverse on the subspace orthogonal to $\mathbf{1}$.

Proof. Since L is symmetric, U is orthogonal and $L = U\Lambda U^T$. The pseudoinverse inverts the nonzero eigenvalues and leaves the zero eigenvalue as zero. Because $\lambda_1 = 0$ corresponds to $u_1 = \mathbf{1}/\sqrt{|V|}$, the sum excludes $i = 1$. $\square$

and its Moore–Penrose pseudoinverse L^+ will govern anatomical friction. The effective resistance between nodes i and j is

$$R_{\text{eff}}(i, j) = (\mathbf{e}_i - \mathbf{e}_j)^T L^+ (\mathbf{e}_i - \mathbf{e}_j) \quad (3)$$

where $\mathbf{e}_i$ is the standard basis vector. This quantity is central to our decision and risk models; it measures how anatomically isolated two locations are, incorporating not just graph distance but the number and capacity of alternative paths.

Proposition 2.2 (Effective Resistance Properties). The effective resistance satisfies: (i) $R_{\text{eff}}(i, j) \geq 0$ with equality iff $i = j$; (ii) symmetry $R_{\text{eff}}(i, j) = R_{\text{eff}}(j, i)$; (iii) triangle inequality $R_{\text{eff}}(i, k) \leq R_{\text{eff}}(i, j) + R_{\text{eff}}(j, k)$; (iv) monotonicity under edge addition: adding an edge cannot increase R_{eff} . Together these imply R_{eff} is a genuine metric on the vertex set, and thus (V_G, R_{eff}) forms a metric space.

Proof sketch. Property (i) follows from positive semi definiteness of L^+ . Property (ii) from symmetry of L^+ . Property (iii) is the triangle inequality for the resistance metric, proved by commuting via the Green’s function interpretation. Property (iv) follows from Rayleigh’s monotonicity principle: adding parallel conductance paths can only decrease effective resistance.

Remark 2.2. The effective resistance has a natural electrical interpretation. If we view G as an electrical network where each edge e_{ij} is a resistor of conductance w_{ij} , then $R_{\text{eff}}(i, j)$ is the electrical resistance measured between nodes i and j when a unit current is injected at i and extracted at j . This interpretation makes it intuitive: well-connected clinics have low effective resistance; a specialty dental clinic reachable only through a single referral corridor has very high effective resistance to the main referral network.

2.2 The HANML Decomposition Principle

For any decision operator $\mathcal{D}$, risk functional $\mathcal{R}$, and stability manifold Ω , the HANML framework decomposes

$$\mathcal{D} = \alpha \mathcal{D}_{\text{AN}} + \beta \mathcal{D}_{\text{NUM}} + \gamma \mathcal{D}_{\text{ML}} \quad (4)$$

with the understanding that this is not a convex combination of outputs but an arbitration over posterior distributions. The weights α, β, γ are not hyperparameters. They are outputs of a Confidence-Weighted Arbitration Protocol (CWAP) defined in Section 3.

Remark 2.3. The additive form (4) is deceptively simple. When $\mathcal{D}_{\text{AN}}$, $\mathcal{D}_{\text{NUM}}$, and $\mathcal{D}_{\text{ML}}$ are posterior distributions (as they are in the MDE synthesis of Section 4.4), the “sum” is actually a log-linear pool: the unnormalized posterior is the product of expert densities raised to powers α, β, γ . This is the Bayesian committee machine of Section 4.4, and it is the sense in which the three modalities “arbitrate” rather than average. The additive notation (4) is the marginal summary of this multiplicative structure.

3. THE CORE HANML ARCHITECTURE

3.1 Triangulated Epistemology

We treat analytical, numerical, and machine learning registers as distinct but non-reducible. The analytical register concerns structure and logical implication. The numerical register concerns temporal evolution and stochastic exploration. The machine learning register concerns empirical pattern and predictive regularity. The HANML claim is that these three registers must be fused through explicit, quantified confidence metrics rather than vague methodological eclecticism.

Definition 3.1 (Epistemic Register). An epistemic register is a pair $(\mathcal{M}, \mathcal{E})$ where $\mathcal{M}$ is a model class and $\mathcal{E}$ is an evidence relation that determines when a model in $\mathcal{M}$ is supported by data. The analytical register has $\mathcal{M}_{\text{AN}} = \text{axiom systems}$ and $\mathcal{E}_{\text{AN}} = \text{logical coherence plus stakeholder consensus}$. The numerical register has $\mathcal{M}_{\text{NUM}} = \text{simulation trajectories}$ and $\mathcal{E}_{\text{NUM}} = \text{convergence plus validation against held-out scenarios}$. The ML register has $\mathcal{M}_{\text{ML}} = \text{trained predictors}$ and $\mathcal{E}_{\text{ML}} = \text{cross-validation accuracy plus domain alignment}$.

Remark 3.1. The non-reducibility claim is important. Analytical models can be wrong about magnitudes but right about structure (e.g., a queueing model that correctly predicts that adding a server reduces waiting time but misestimates the magnitude). Numerical simulations can be calibrated to magnitudes but obscure structural assumptions (e.g., a Monte Carlo that matches a Sharpe ratio but hides its assumption of normal returns). ML predictors can match patterns but violate structural constraints (e.g., a neural network that predicts demand without respecting conservation laws). No single register subsumes the others; the HANML framework requires all three.

3.2 The Confidence-Weighted Arbitration Protocol

Let $C_{AN}(t)$, $C_{NUM}(t)$, and $C_{ML}(t)$ denote time-varying epistemic confidence scores for each register. The CWAP assigns arbitration weights via a temperature-scaled softmax:

$$\alpha(t) = \frac{\exp(\kappa C_{AN}(t))}{Z(t)}, \quad \beta(t) = \frac{\exp(\kappa C_{NUM}(t))}{Z(t)}, \quad \gamma(t) = \frac{\exp(\kappa C_{ML}(t))}{Z(t)} \quad (5)$$

where

$$Z(t) = \exp(\kappa C_{AN}(t)) + \exp(\kappa C_{NUM}(t)) + \exp(\kappa C_{ML}(t)) \quad (6)$$

and $\kappa > 0$ is an inverse-temperature parameter. In the limit $\kappa \rightarrow \infty$, the framework collapses to the single most confident modality. In the limit $\kappa \rightarrow 0$, it averages blindly. We typically set $\kappa \approx 2$ to allow meaningful arbitration without winner-take-all pathology.

Derivation of the Softmax Form. The softmax form (5) is not arbitrary; it emerges from a maximum-entropy principle. Suppose we seek the least-informative distribution (α, β, γ) on the simplex that respects the constraint that the expected confidence equals a target $\bar{C}$:

$$\alpha C_{AN} + \beta C_{NUM} + \gamma C_{ML} = \bar{C}, \quad \alpha + \beta + \gamma = 1, \quad \alpha, \beta, \gamma \geq 0$$

Maximizing the Shannon entropy $H = -\sum_w w \log w$ subject to these constraints via Lagrange multipliers yields

$$w_i = \frac{\exp(\lambda C_i)}{\sum_j \exp(\lambda C_j)}$$

with λ the dual variable on the confidence constraint. Setting $\lambda = \kappa$ recovers (5). Thus the CWAP weights are the maximum-entropy distribution consistent with the observed confidence scores, which is the least committal distribution the framework can adopt without additional information.

Proposition 3.1 (Limiting Behaviour). (i) As $\kappa \rightarrow \infty$, the weights concentrate on the modality with maximal $C_*(t)$. (ii) As $\kappa \rightarrow 0^+$, the weights converge to the uniform distribution $(1/3, 1/3, 1/3)$. (iii) The weights are differentiable in κ for $\kappa > 0$, and the derivative is bounded by $\kappa \cdot \text{Var}(C)$.

Proof. (i) By Laplace's method, $\exp(\kappa C_i) / \sum_j \exp(\kappa C_j) \rightarrow 1_{[i = \arg\max_j C_j]}$ as $\kappa \rightarrow \infty$. (ii) Direct: $\exp(0) / (3\exp(0)) = 1/3$. (iii) Differentiate the softmax; the Hessian is the covariance of the indicator under the softmax distribution, bounded by $1/4$. $\square$

The confidence functions are constructed as follows. Analytical confidence aggregates theoretical coherence, axiomatic completeness, and stakeholder consensus:

$$C_{AN}(t) = \sigma_1 \rho_{\text{coherence}}(t) + \sigma_2 \rho_{\text{axiomatic}}(t) + \sigma_3 \rho_{\text{consensus}}(t) \quad (7)$$

with $\sigma_1 + \sigma_2 + \sigma_3 = 1$. Numerical confidence reflects simulation quality:

$$C_{NUM}(t) = \delta_1 \rho_{\text{convergence}}(t) + \delta_2 (1 - S_{\text{max}}(t)) + \delta_3 \rho_{\text{validation}}(t) \quad (8)$$

where $S_{\text{max}}(t)$ is the maximum sensitivity index of simulation outputs to input parameters, so that $(1 - S_{\text{max}})$ measures robustness. Machine learning confidence tracks predictive integrity:

$$C_{ML}(t) = \eta_1 A_{CV}(t) + \eta_2 (1 - V_{\text{feat}}(t)) + \eta_3 \rho_{\text{domain}}(t) \quad (9)$$

where A_{CV} is stratified cross-validation accuracy, V_{feat} is the variance of feature importance rankings across bootstrap samples, and ρ_{domain} measures alignment with established domain knowledge.

Remark 3.2 (Construction of the Sub-Scores). The sub-scores $\rho_{\text{coherence}}$, $\rho_{\text{axiomatic}}$, $\rho_{\text{consensus}}$ are each normalized to $[0, 1]$. Coherence is operationalized as $1 - \|f(x^*) - f(x)\| / \|f(x^*)\|$ where x^* is the equilibrium implied by axioms and x is the current operational state; small residuals mean high coherence. Axiomatic completeness is the fraction of Arrow-style axioms (completeness, transitivity, independence of irrelevant alternatives, non-dictatorship) that the decision problem satisfies. Stakeholder consensus is measured via a Delphi-style elicitation: $\rho_{\text{consensus}} = 1 - \text{Var}_i[\hat{v}_i] / \text{Var}_i[\hat{v}_i]_{\text{max}}$ where $\hat{v}_i$ is stakeholder i 's valuation. The weights $\sigma_1, \sigma_2, \sigma_3$ reflect clinical context; in heavily pluralistic settings (e.g., state-board disputes), σ_3 is set high.




Figure 1: The HANML Triangulated Architecture and CWAP Arbitration Flow.

4. THE META-DECISION ENGINE (MDE)

Decision-making in dental practice is rarely a single choice. It is an iterative, spatially distributed, partially observable process. The MDE fuses three mathematical layers into a hierarchical Bayesian committee.

4.1 Analytical Layer: Culturally Weighted TOPSIS

Consider a set of alternatives $\{a_1, \dots, a_k\}$ and an expanded criteria set $C = \{c_1, \dots, c_n\} \cup \{c_{n+1}, \dots, c_{n+q}\}$. The additional criteria capture clinical dimensions: alignment with national dental association targets (c_{formal}), dental-society priority alignment (c_{informal}), reimbursement impact on staff retention (c_{remit}), and anatomical accessibility improvement measured via referral-graph centrality (c_{topo}).

Because several criteria resist crisp measurement—notably c_{informal} —we employ triangular fuzzy numbers. The rating of alternative a_i on criterion c_j is $\tilde{x}_{ij} = (l_{ij}, m_{ij}, u_{ij})$. The fuzzy decision matrix is $\tilde{X} = [\tilde{x}_{ij}]$.

Definition 4.1 (Triangular Fuzzy Number). A triangular fuzzy number $\tilde{a} = (l, m, u)$ has membership function

$$\mu_{\tilde{a}}(x) = \begin{cases} 0 & x < l \text{ or } x > u, \\ (x - l)/(m - l) & l \leq x \leq m, \\ (u - x)/(u - m) & m \leq x \leq u. \end{cases}$$

The mode m is the most likely value, l the optimistic lower bound, and u the pessimistic upper bound. Fuzzy arithmetic is performed componentwise: $\tilde{a} + \tilde{b} = (l_a + l_b, m_a + m_b, u_a + u_b)$ and $\tilde{a} \cdot \tilde{b} \approx (l_a l_b, m_a m_b, u_a u_b)$ for positive fuzzy numbers.

The weight vector is not exogenous. It emerges from a cooperative game among stakeholder groups $N = \{1, \dots, s\}$. The characteristic function $v: 2^N \rightarrow \mathbb{R}_+$ assigns to each coalition $S \subseteq N$ the maximum TOPSIS score achievable if that coalition unilaterally set weights. The Shapley value for stakeholder i is

$$\phi_i(v) = \sum_{S \subseteq N \setminus \{i\}} \frac{|N|!}{|S|! (|N| - |S| - 1)!} [v(S \cup \{i\}) - v(S)] \quad (10)$$

Derivation of the Shapley Formula. The Shapley value is the unique allocation satisfying four axioms: efficiency ($\sum_i \phi_i = v(N)$), symmetry (interchangeable players get equal shares), dummy player (players who never contribute get zero), and additivity ($\phi_i(v + w) = \phi_i(v) + \phi_i(w)$). To derive the formula, consider adding player i to a random permutation of $N \setminus \{i\}$. The set of players preceding i in a uniformly random permutation is uniformly distributed over all $S \subseteq N \setminus \{i\}$, with probability $|S|! (|N| - |S| - 1)! / |N|!$ of being exactly S . Player i 's marginal contribution when joining after S is $v(S \cup \{i\}) - v(S)$. Averaging over all permutations yields (10).

The compromise weight for criterion j is then

$$w_j = \sum_{i \in N} \phi_i(v) w_j^{(i)} \quad (11)$$

where $w_j^{(i)}$ is player i 's ideal weight. This ensures that a general dentist and a specialist planner, or a regulatory official and a local dental-society elder, exert influence proportional to their coalitional leverage rather than procedural dominance.

Proposition 4.1 (Efficiency of Compromise Weights). If each player's ideal weight vector $w^{(i)}$ sums to 1, then the compromise weight vector w also sums to 1.

Proof. $\sum_j w_j = \sum_j \sum_i \phi_i(v) w_j^{(i)} = \sum_i \phi_i(v) \sum_j w_j^{(i)} = \sum_i \phi_i(v) \cdot 1 = v(N)/v(N) = 1$, using efficiency of the Shapley value and the assumption $\sum_j w_j^{(i)} = 1$. $\square$

The fuzzy positive-ideal and negative-ideal solutions are $\tilde{A}^+ = (\tilde{v}_1^+, \dots, \tilde{v}_{n+q}^+)$ and $\tilde{A}^- = (\tilde{v}_1^-, \dots, \tilde{v}_{n+q}^-)$, where $\tilde{v}_j^+ = (1, 1, 1)$ for benefit criteria and $\tilde{v}_j^- = (0, 0, 0)$. The separation measures use the vertex method:

$$\tilde{d}_{ij}^+ = \frac{1}{3} \sqrt{[(l_{ij} - l_j^+)^2 + (m_{ij} - m_j^+)^2 + (u_{ij} - u_j^+)^2]} \quad (12)$$

with $\tilde{d}_{ij}^-$ defined analogously. The factor $1/3$ averages over the three vertices of the triangular fuzzy number, providing a crisp distance.

Remark 4.1 (Choice of Distance). The vertex method (12) is one of several fuzzy distance metrics. Alternatives include the α -cut integral distance and the Hausdorff distance. We use the vertex method because it is computationally trivial and preserves the interpretability of “ideal” and “anti-ideal” anchors. For dental clinical applications where stakeholders must verify the calculations by hand, transparency outweighs metric sophistication.

The relative closeness to the ideal is

$$\tilde{C}_i = \frac{\sum_j w_j \tilde{d}_{ij}^-}{\sum_j w_j \tilde{d}_{ij}^- + \sum_j w_j \tilde{d}_{ij}^+} \quad (13)$$

and the analytical decision potential is $\phi_{AN}(a_i) = \tilde{C}_i$.

Proposition 4.2 (Range of Closeness). The closeness $\tilde{C}_i \in [0,1]$, with $\tilde{C}_i = 1$ iff alternative a_i coincides with the positive ideal $\tilde{A}^+$ on all criteria, and $\tilde{C}_i = 0$ iff a_i coincides with the negative ideal $\tilde{A}^-$.

Proof. If $\tilde{d}_{ij}^+ = 0$ for all j , then $\tilde{C}_i = \sum_j w_j \tilde{d}_{ij}^- / (\sum_j w_j \tilde{d}_{ij}^- + 0) = 1$. Conversely, if $\tilde{d}_{ij}^- = 0$ for all j , then $\tilde{C}_i = 0 / (0 + \sum_j w_j \tilde{d}_{ij}^+) = 0$. The intermediate values follow from continuity and monotonicity in the distances. $\square$

4.2 Numerical Layer: Anatomically Constrained POMDPs

Decisions propagate through dental practice’s clinical variability imperfectly. We model this as a Partially Observable Markov Decision Process (POMDP) with anatomically modulated transitions. The standard POMDP is $(S, A, T, R, O, Z, \gamma_{disc})$. We modify the transition kernel by an anatomical friction function:

$$T_G(s' | s, a, \theta) = T(s' | s, a) \phi_G(\text{loc}(s), \text{loc}(s'), \theta) \psi(\text{season}(t)) \quad (14)$$

where $\text{loc}(s) \in V_G$ extracts the location component, ψ is a seasonal modulation with period $T = 12$ months, and

$$\phi_G(v_i, v_j) = \frac{1}{1 + R_{eff}(v_i, v_j)} \quad (15)$$

using the effective resistance (3). Actions between well-connected locations propagate reliably; actions across anatomically isolated specialty-boundary zones or seasonal patient flow-severed specialty-practice corridors attenuate.

Derivation of the Friction Function. The form (15) is chosen so that $\phi_G(v_i, v_i) = 1$ (no friction at the same location) and $\phi_G(v_i, v_j) \rightarrow 0$ as $R_{eff}(v_i, v_j) \rightarrow \infty$ (complete attenuation for infinitely isolated nodes). The reciprocal form ensures the friction is bounded: even maximally isolated nodes retain a small positive probability of influence, which is the correct asymptotic for rare but possible communication (e.g., helicopter supply). The function is differentiable in R_{eff} , which permits gradient-based optimization in the policy iteration step.

The belief state $b_t(s) = P(s_t = s | o_{1:t}, a_{1:t-1})$ updates via the standard Bayesian filter but with the modified kernel:

$$b_{t+1}(s') = \frac{Z(o_{t+1} | s', a_t) \sum_s T_G(s' | s, a_t, \theta_t) b_t(s)}{\sum_{s''} Z(o_{t+1} | s'', a_t) \sum_s T_G(s'' | s, a_t, \theta_t) b_t(s)} \quad (16)$$

Derivation of the Belief Update. Starting from the definition $b_{t+1}(s') = P(s_{t+1} = s' | o_{1:t+1}, a_{1:t})$, apply Bayes’ rule:

$$b_{t+1}(s') \propto P(o_{t+1} | s', a_{1:t}, o_{1:t}) P(s_{t+1} = s' | a_{1:t}, o_{1:t})$$

The observation likelihood factorizes as $Z(o_{t+1} | s', a_t)$ by the Markov assumption on observations. The predictive prior is obtained by marginalizing over s_t :

$$P(s_{t+1} = s' | a_{1:t}, o_{1:t}) = \sum_s T_G(s' | s, a_t, \theta_t) b_t(s)$$

Combining yields the numerator of (16); the denominator normalizes by the marginal likelihood of o_{t+1} .

Exact solution is intractable. We employ point-based value iteration, approximating the value function by a finite set of α -vectors Γ :

$$V(b) = \max_{\alpha \in \Gamma} \sum_s \alpha(s) b(s) \quad (17)$$

Remark 4.2 (Point-Based Value Iteration). The α -vector representation exploits the piecewise-linearity and convexity of the POMDP value function. Each $\alpha \in \Gamma$ corresponds to a policy tree, and $\alpha(s)$ is the expected discounted reward of executing that tree starting from state s . The maximization over Γ selects the best policy tree for the current belief. Point-based methods sample a finite set of beliefs $\mathcal{B}$ (typically reachable from the initial belief) and update Γ only on $\mathcal{B}$, reducing the exponential complexity of exact value iteration to polynomial in $|\mathcal{B}|$.

The numerical decision potential $\phi_{NUM}(a)$ is the expected α -vector value of executing action a from the current belief.

4.3 Machine Learning Layer: SDODL

The Sparse Dental Operational Data Learner (SDODL) addresses the small-N, high-D pathology. It has three modules. **Transfer learning with domain adaptation.** Given source domains D_S (ambulatory clinic data, specialty practice data) and target D_T (dental-clinical data), we minimize

$$\min_{\theta_f} W_l(P_S^{\theta_f}, P_T^{\theta_f}) + \lambda L_{task}(\theta_f, \theta_y; D_T) \quad (18)$$

where the 1-Wasserstein distance is computed via the Kantorovich–Rubinstein dual:

$$W_1(P, Q) = \sup_{\|\eta\|_L \leq 1} \mathbb{E}_{x \sim P}[f(x)] - \mathbb{E}_{x \sim Q}[f(x)] \quad (19)$$

Derivation of the Kantorovich–Rubinstein Dual. The primal 1-Wasserstein problem is

$$W_1(P, Q) = \inf_{\pi \in \Pi(P, Q)} \mathbb{E}_{(x, y) \sim \pi} [\|x - y\|]$$

where $\Pi(P, Q)$ is the set of joint distributions with marginals P and Q . Lagrangian duality yields

$$W_1(P, Q) = \sup_f \{ \mathbb{E}_{x \sim P}[f(x)] - \mathbb{E}_{y \sim Q}[f(y)] : f \text{ is 1-Lipschitz} \}$$

The Lipschitz constraint $\|f\|_L \leq 1$ arises because the dual variable on the marginal constraints is precisely the test function f , and the cost $\|x - y\|$ is 1-Lipschitz. Enforcing the Lipschitz constraint via gradient penalty (as in WGAN-GP) gives a tractable training objective.

and the Lipschitz constraint is enforced by gradient penalty.

Relational graph convolution for informal professional networks. The heterogeneous graph H has node types $\tau: V_H \rightarrow \mathcal{T}$ and relation types $\phi: E_H \rightarrow \mathcal{R}$. The R-GCN propagation is

$$h_i^{(l+1)} = \sigma \left(\sum_{r \in \mathcal{R}} \sum_{j \in N_i^r} \frac{1}{c_{i,r}} W_r^{(l)} h_j^{(l)} + W_0^{(l)} h_i^{(l)} \right) \quad (20)$$

where $c_{i,r}$ is a normalization constant and $W_r^{(l)}$ are relation-specific weights. Fuzzy edge memberships $\mu_{ij} \in [0, 1]$ replace binary adjacency, so the sum is weighted by μ_{ij} .

Remark 4.3 (Fuzzy Edge Weighting). Replacing binary adjacency with fuzzy memberships μ_{ij} allows the R-GCN to capture graded relationships: a dental society elder's influence on a regulatory official is not binary (present/absent) but graded by frequency of contact, referral proximity, and professional obligation. The fuzzy memberships are elicited via ethnographic survey and normalized so that $\sum_j \mu_{ij} = 1$ for each node, preserving the convolution interpretation.

Temporal fusion with festival encoding. We use a Temporal Fusion Transformer (TFT). The multi-horizon attention is

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} + M \right) V \quad (21)$$

where M masks known future inputs (seasonality, festival calendars) from unknown stochastic shocks. Festival indicators $f_k(t) \in \{0, 1\}$ and seasonal case-mix phase $m(t) \in \{0, 1, 2\}$ are embedded and added to the static covariate representation. Derivation of the Masked Attention. Standard scaled dot-product attention computes $\text{softmax}(QK^T / \sqrt{d_k})V$. The scaling by $1/\sqrt{d_k}$ prevents the dot products from growing large in high dimensions, which would push the softmax into saturation. The mask M is a matrix with $-\infty$ in positions corresponding to disallowed (future, unknown) tokens and 0 elsewhere; adding M before the softmax forces the attention weights on masked positions to zero. In the TFT, the mask ensures that predictions at horizon $t + h$ attend only to information available at time t , preventing leakage.

The ML decision potential $\phi_{ML}(a)$ is the predicted probability of successful outcome under action a , derived from the TFT output head.

4.4 Synthesis: The Bayesian Committee Machine

The MDE combines the three potentials through a hierarchical Bayesian committee. The posterior predictive over outcome y given action a , state x , and parameters θ is

$$P(y | a, x, \theta) \propto P_{AN}(y | a, x, \theta)^\alpha P_{NUM}(y | a, x, \theta)^\beta P_{ML}(y | a, x, \theta)^\gamma \quad (22)$$

where the exponents are the CWA weights from (5). This is not a simple product of experts; the exponents perform confidence-scaled tempering of each posterior.

Derivation of the Log-Linear Pool. A logarithmic opinion pool of K experts with weights w_k is

$$P_{\text{pool}}(y) \propto \prod_{k=1}^K P_k(y)^{w_k}$$

Taking logs, $\log P_{\text{pool}}(y) = \sum_k w_k \log P_k(y) + \text{const}$, which is a weighted average of log-densities. When $w_k = 1/K$, this reduces to the geometric mean of densities. The exponents α, β, γ in (22) are the CWA weights, so high-confidence modalities contribute more logarithmically. The pool is proper (integrates to 1 after normalization) and externally Bayesian (combining priors then updating equals updating then combining). These properties make it the canonical choice for multi-expert fusion.

Proposition 4.3 (Properness of the Pool). The log-linear pool (22) is a proper probability density after normalization, provided each $P_\star$ is proper and at least one weight is positive.

Proof. $\int P_{AN}^\alpha P_{NUM}^\beta P_{ML}^\gamma dy \leq (\int P_{AN} dy)^\alpha (\int P_{NUM} dy)^\beta (\int P_{ML} dy)^\gamma = 1$ by Hölder's inequality, with equality iff the densities are proportional. Thus the integral is finite and positive, and normalization yields a proper density. $\square$

Rather than a single optimal action, the MDE outputs a decision robustness profile $\mathcal{P}(a)$:

$$\mathcal{P}(a) = (\mathbb{E}[y | a], \text{Var}[y | a], S_{AN}(a), S_{NUM}(a), S_{ML}(a)) \quad (23)$$

where $S_\star(a) = \|\nabla_\star \mathbb{E}[y | a]\|$ measures sensitivity to assumptions in modality $\star$. The recommended action satisfies robust satisficing:

$$a^* = \arg \min_{a: \mathbb{E}[y|a] \geq \tau} [\lambda_{AN} S_{AN}(a) + \lambda_{NUM} S_{NUM}(a) + \lambda_{ML} S_{ML}(a)] \quad (24)$$

with aspiration level τ and modality penalty weights λ .

Remark 4.4 (Satisficing vs. Optimizing). The robust satisficing criterion (24) is deliberately not expected-utility maximization. Under deep uncertainty (Section 7), expected utility is fragile because the expectation itself depends on contested probability assignments. Satisficing—meeting an aspiration level while minimizing sensitivity—is more robust. The aspiration τ is set by organizational leadership: “we need at least τ expected successful days per month.” The penalty weights λ_* reflect which modalities’ assumptions are most contested; in a data-poor environment, λ_{ML} is set high so that ML-sensitive actions are penalized.




Figure 2: The Meta-Decision Engine (MDE) Computational Graph.

5. THE DYNAMIC RISK RESILIENCE MODEL (DRRM)

5.1 Clinical Fault Trees with Imprecise Probability

Classical fault trees fail in dental practice because they omit clinical failure modes. We extend the basic event set to include B_1 : state-board jurisdictional conflict; B_2 : dental-society opposition; B_3 : reimbursement corridor disruption; B_4 : seasonal schedule disruption exceeding 14 days.

Definition 5.1 (Fault Tree Structure Function). A fault tree on basic events $B_1, \dots, B_n$ is a Boolean function $g: \{0,1\}^n \rightarrow \{0,1\}$ defined by AND and OR gates. The system fails ($g = 1$) when the gate logic is satisfied. For coherent systems, g is monotone non-decreasing in each argument and $g(0, \dots, 0) = 0$, $g(1, \dots, 1) = 1$.

Because historical frequencies for these events are scarce, we represent basic event probabilities as probability boxes (p-boxes):

$$\underline{F}_i(x) \leq F_i(x) \leq \bar{F}_i(x) \quad \forall x \quad (25)$$

The system failure probability is interval-valued:

$$\underline{P}_{\text{sys}} = g(\underline{F}_1, \dots, \underline{F}_n), \quad \bar{P}_{\text{sys}} = g(\bar{F}_1, \dots, \bar{F}_n) \quad (26)$$

where g is the fault tree structure function. For coherent systems, these bounds are computed via interval arithmetic on the minimal cut sets.

Derivation of the P-Box Bounds. A p-box $[\underline{F}, \bar{F}]$ represents all cumulative distribution functions F satisfying $\underline{F}(x) \leq F(x) \leq \bar{F}(x)$. For a coherent fault tree (monotone structure function), the system failure probability is monotone in each basic event probability. Therefore:

The lower bound $\underline{P}_{\text{sys}}$ is achieved by setting each F_i to its lower bound $\underline{F}_i$ and propagating through g via interval arithmetic.

The upper bound $\bar{P}_{\text{sys}}$ is achieved by setting each F_i to its upper bound $\bar{F}_i$.

For AND gates (parallel failure: system fails only if all events occur), $P_{\text{AND}} = \prod_i P_i$, so the bounds multiply. For OR gates (series failure: system fails if any event occurs), $P_{\text{OR}} = 1 - \prod_i (1 - P_i)$, so the bounds combine via complements. Minimal cut sets decompose any coherent fault tree into a disjunction of AND gates, enabling systematic computation.

Remark 5.1 (Why P-Boxes for Dental Practice). Classical fault trees assume point probabilities for basic events. In dental practice, the historical record for clinical events like B_1 (jurisdictional conflict) is too short to support point estimates; the current dental regulatory framework is less than a decade old. P-boxes encode epistemic uncertainty directly in the probability specification, propagating it through the fault tree without forcing false precision. The resulting interval $[\underline{P}_{\text{sys}}, \bar{P}_{\text{sys}}]$ is honest about what we do not know.

5.2 Agent-Based Contagion and the Specialty-Boundary Effect

Risk cascades in dental practice do not diffuse isotropically. They channel along specialty-practice corridors and referral routes. We model this through an agent-based framework combined with anisotropic diffusion.

Each agent i carries state $s_i(t) \in \mathbb{R}^d$ (financial health, operational capacity, referral network density, anatomical coordinates). The dynamics are

$$\frac{ds_i}{dt} = f(s_i, N_i, \theta) + \xi_i(t) + J_i(t) \quad (27)$$

where N_i is the neighborhood determined by anatomical and professional proximity, $\xi_i(t)$ is Gaussian noise, and $J_i(t)$ is a compound Poisson jump process capturing discrete shocks:

$$J_i(t) = \sum_{k=1}^{N_i} Y_k \delta(t - T_k) \quad (28)$$

with jump marks Y_k drawn from a heavy-tailed distribution (Lévy α -stable) to reflect procedural-complication and postoperative-sequela severity.

Remark 5.2 (Choice of Lévy-Stable Jumps). Gaussian jumps would underestimate tail risk because procedural complication magnitudes follow a power law (Gutenberg–Richter). Lévy α -stable distributions with $\alpha \in (1, 2)$ have power-law tails $\sim |Y|^{-(1+\alpha)}$ and are the attractor of generalized central limit theorems for heavy-tailed sums. Setting $\alpha \approx 1.5$ matches empirical complication severity distributions. The cost is that α -stable distributions with $\alpha < 2$ have infinite variance, which complicates estimation but correctly reflects the unbounded nature of catastrophic risk.

The anatomical risk diffusion tensor $T(x, y)$ modifies standard isotropic diffusion. The risk density $\rho(x, y, t)$ evolves according to

$$\frac{\partial \rho}{\partial t} = \nabla \cdot (T(x, y) \nabla \rho) + S(x, y, t) - \mu \rho \quad (29)$$

where S is the source term and μ the recovery rate. The tensor is diagonalized as $T = R^T \Lambda R$, where R rotates coordinates to align with local referral-corridor directions and $\Lambda = \text{diag}(\lambda_{\max}, \lambda_{\min})$ with $\lambda_{\max} \gg \lambda_{\min}$. This is the Specialty-Boundary Effect: risk elongates along anatomical gradients.

Derivation of the Anisotropic Diffusion Equation. Start from Fick's first law of diffusion, which states that the flux J of risk is proportional to the negative gradient of risk density: $J = -T \nabla \rho$, where T is the diffusion tensor. Conservation of risk mass gives the continuity equation $\partial \rho / \partial t = -\nabla \cdot J + S - \mu \rho$, where S is the source and $\mu \rho$ the decay. Substituting Fick's law:

$$\frac{\partial \rho}{\partial t} = \nabla \cdot (T \nabla \rho) + S - \mu \rho$$

In isotropic diffusion, $T = D I$ for scalar diffusivity D , recovering Fick's second law. In our model, T is spatially varying and anisotropic: $\lambda_{\max}$ aligns with the local referral-corridor direction (high diffusivity along referral corridors and supply lines) and $\lambda_{\min}$ is orthogonal (low diffusivity across specialty-boundary zones). The rotation matrix $R(x, y)$ is constructed from the local anatomical gradient, computed from a digital elevation model. The result is that risk spreads rapidly along the main referral corridor but slowly across the specialty-boundary zone—a faithful representation of the Specialty-Boundary Effect.

5.3 Graph Neural Network Risk Clustering

We construct a risk association graph $R = (V_R, E_R)$ where nodes are risk events and edges are weighted by historical co-occurrence. A Graph Attention Network (GAT) learns embeddings:

$$h_i' = \sigma \left(\sum_{j \in N_i} \alpha_{ij} W h_j \right) \quad (30)$$

with attention coefficients

$$e_{ij} = \text{LeakyReLU}(a^T [W h_i \parallel W h_j]), \quad \alpha_{ij} = \frac{\exp(e_{ij})}{\sum_{k \in N_i} \exp(e_{ik})} \quad (31)$$

Derivation of the GAT Attention. The attention mechanism computes a relevance score e_{ij} for each neighbor j of node i . The score is computed by applying a shared linear transformation W to both h_i and h_j , concatenating, applying a learnable attention vector a , and passing through a LeakyReLU nonlinearity. The softmax over neighbors normalizes the attention coefficients to a probability distribution. The updated embedding h_i' is a weighted average of transformed neighbor embeddings, with weights given by the attention coefficients. Multi-head attention runs K independent attention heads in parallel and concatenates or averages their outputs, stabilizing training and allowing the model to attend to different aspects of the neighborhood simultaneously.

The GNN predicts emergent risk clusters—latent combinations of events with high structural similarity but no direct historical co-occurrence.

5.4 The Risk Resilience Index

For organization o , the DRRM outputs

$$\text{RRI}_o = \frac{\omega_{\text{abs}} R_{\text{absorptive}}^{(\text{AN})} + \omega_{\text{adt}} R_{\text{adaptive}}^{(\text{NUM})} + \omega_{\text{ant}} R_{\text{anticipatory}}^{(\text{ML})}}{R_{\text{total}}} \quad (32)$$

where

$$R_{\text{absorptive}}^{(\text{AN})} = 1 - \bar{P}_{\text{sys}}, \quad R_{\text{adaptive}}^{(\text{NUM})} = \mathbb{E}[\tau_{\text{rec}} | \text{shock}]^{-1}, \quad R_{\text{anticipatory}}^{(\text{ML})} = \text{Precision}_{\text{GNN}} \quad (33)$$

and τ_{rec} is recovery time from ABM simulations. The weights ω reflect strategic posture.

Proposition 5.1 (Decomposition of Resilience). The three resilience components correspond to the three temporal phases of risk response: absorptive capacity is the ability to withstand the initial shock (analytical, structural), adaptive capacity is the ability to recover during the aftermath (numerical, dynamic), and anticipatory capacity is the ability to foresee and prepare (ML, predictive). The HANML arbitration ensures each phase is handled by the modality best suited to its time scale.

Proof. The decomposition follows from the time scales: analytical fault trees capture static structural capacity (pre-shock), ABM simulations capture dynamic recovery (post-shock), and GNN predictions capture future risk patterns (forward-looking). The weights ω encode the organization's strategic posture: a risk-averse organization sets ω_{abs} high; a proactive organization sets ω_{ant} high. $\square$




Figure 3: The Dynamic Risk Resilience Model (DRRM) and Specialty-Boundary Effect.

6. THE STOCHASTIC STABILITY SURFACE (SSS)

6.1 Three Metrics of Viability

Operational stability is not variance minimization. It is persistence within a viability kernel. We define three metrics. Structural stability derives from the analytical Jacobian of resource flows. Let $\dot{x} = f_{\text{AN}}(x)$ describe the unperturbed system with equilibrium x^* . The Jacobian $J = \nabla f_{\text{AN}}|_{x^*}$ yields eigenvalues $\lambda_i(J)$. Then

$$S_{\text{struct}}(x^*) = -\max_i \text{Re}(\lambda_i(J)) \quad (34)$$

Positive S_{struct} indicates local exponential stability.

Derivation. Linearize $\dot{x} = f_{\text{AN}}(x)$ around x^* : $\delta\dot{x} = J\delta x + O(\|\delta x\|^2)$, where $\delta x = x - x^*$. The solution is $\delta x(t) = e^{Jt}\delta x(0)$. Decomposing into eigenmodes, $\delta x(t) = \sum_i c_i e^{\lambda_i t} v_i$ where λ_i are eigenvalues and v_i eigenvectors. The mode i decays iff $\text{Re}(\lambda_i) < 0$. The slowest-decaying mode dominates long-term behavior, with rate $\max_i \text{Re}(\lambda_i)$. Therefore local exponential stability requires $\max_i \text{Re}(\lambda_i) < 0$, and the stability margin is $-\max_i \text{Re}(\lambda_i) = S_{\text{struct}}$.

Dynamical stability captures stochastic perturbations. The organization evolves as a jump-diffusion:

$$dx_t = f_{\text{NUM}}(x_t) dt + \Sigma(x_t) dW_t + dL_t \quad (35)$$

where W_t is standard Brownian motion and L_t the compound Poisson process (28). The largest Lyapunov exponent is

$$\lambda_{\text{max}} = \lim_{t \rightarrow \infty} \frac{1}{t} \mathbb{E} \left[\ln \frac{\|\delta x(t)\|}{\|\delta x(0)\|} \right] \quad (36)$$

Derivation of the Lyapunov Exponent. Consider two nearby trajectories $x(t)$ and $x(t) + \delta x(t)$ of the stochastic system (35). The separation $\delta x(t)$ evolves (to leading order) by the variational equation $d(\delta x) = \nabla f_{\text{NUM}}(x_t) \delta x dt + \nabla \Sigma(x_t) \delta x dW_t$. The Lyapunov exponent measures the asymptotic exponential rate of separation:

$$\lambda_{\text{max}} = \lim_{t \rightarrow \infty} \frac{1}{t} \ln \frac{\|\delta x(t)\|}{\|\delta x(0)\|}$$

For stochastic systems, we take the expectation to handle the noise. If $\lambda_{\text{max}} < 0$ almost surely, nearby trajectories converge and the system is dynamically stable. If $\lambda_{\text{max}} > 0$, nearby trajectories diverge and the system is chaotic or unstable.

Numerically, we estimate via Wolf's algorithm:

$$\lambda_{\text{max}} \approx \frac{1}{N \Delta t} \sum_{k=1}^N \ln \frac{\|\delta x(t_k + \Delta t)\|}{\|\delta x(t_k)\|} \quad (37)$$

with periodic renormalization of the separation vector δx . Dynamical stability requires $\lambda_{\text{max}} < 0$ almost surely.

Remark 6.1 (Wolf's Algorithm). Wolf's algorithm tracks a reference trajectory and a perturbed trajectory. At each time step, the perturbed trajectory is renormalized to a fixed separation $\|\delta x\| = \epsilon$ after the growth factor is logged. This prevents numerical overflow and ensures the linearization remains valid. The sum of log-growth factors, divided by total time, converges to the largest Lyapunov exponent under ergodicity assumptions.

Predictive stability uses a Variational Autoencoder. The encoder $q_\phi(z|x)$ and decoder $p_\theta(x|z)$ are trained to maximize the ELBO:

$$\mathcal{L}_{\text{ELBO}}(\phi, \theta; x) = \mathbb{E}_{q_\phi(z|x)} [\log p_\theta(x|z)] - \beta \text{KL}(q_\phi(z|x) \parallel p(z)) \quad (38)$$

with $\beta > 1$ for disentangled representations. Predictive stability is

$$S_{\text{pred}}(x) = \mathcal{L}_{\text{ELBO}}(\phi, \theta; x) \quad (39)$$

Derivation of the ELBO. The evidence lower bound derives from Jensen's inequality applied to the log-marginal likelihood:

$$\log p(x) = \log \int p(x, z) dz = \log \int p(x|z) p(z) dz = \log \mathbb{E}_{q_\phi(z|x)} \left[\frac{p(x|z) p(z)}{q_\phi(z|x)} \right]$$

By Jensen's inequality (log is concave):

$$\log p(x) \geq \mathbb{E}_{q_\phi(z|x)} \left[\log \frac{p(x|z) p(z)}{q_\phi(z|x)} \right] = \mathbb{E}_{q_\phi} [\log p(x|z)] - \text{KL}(q_\phi(z|x) \parallel p(z))$$

This is the standard ELBO. The β -VAE modification scales the KL term by $\beta > 1$, which encourages disentangled latent representations by tightening the prior constraint. Low ELBO indicates that the input x is poorly reconstructed by the VAE, signaling an anomalous state outside the training distribution—a precursor to instability.

Low ELBO signals anomalous, potentially precarious state.

6.2 The Stability Manifold

The Stochastic Stability Surface lives in $(S_{\text{struct}}, S_{\text{dyn}}, S_{\text{pred}})$ -space. The viability region is

$$\mathcal{V} = \{x \in X \mid S_{\text{struct}} \geq \tau_{\text{struct}}, S_{\text{dyn}} \geq \tau_{\text{dyn}}, S_{\text{pred}} \geq \tau_{\text{pred}}\} \quad (40)$$

Remark 6.2 (Viability Theory). The viability region $\mathcal{V}$ is inspired by Aubin's viability theory, which defines the viability kernel as the set of states from which there exists a control keeping the system within prescribed constraints indefinitely. Our definition (40) is a static approximation: we require all three-stability metrics to exceed thresholds simultaneously. The dynamic viability kernel would require solving a Hamilton-Jacobi-Bellman equation, which is intractable in high dimensions. The static approximation is conservative but tractable.

dental-clinical organizations typically operate closer to the boundary $\partial\mathcal{V}$ than their larger corporate counterparts. However, they possess informal elasticity—referral-based credit, on-call staff pool labor reserves—that acts as a fuzzy buffer. We model this by a membership function

$$\mu_{\mathcal{V}}(x) = \min \left\{ \sigma \left(\frac{S_{\text{struct}} - \tau_{\text{struct}}}{\delta_{\text{struct}}} \right), \sigma \left(\frac{S_{\text{dyn}} - \tau_{\text{dyn}}}{\delta_{\text{dyn}}} \right), \sigma \left(\frac{S_{\text{pred}} - \tau_{\text{pred}}}{\delta_{\text{pred}}} \right) \right\} \quad (41)$$

where σ is the logistic function. The fuzzy boundary captures that an organization may violate formal thresholds yet remain viable through community support.

Proposition 6.1 (Fuzzy Viability). The membership function $\mu_{\mathcal{V}}(x) \in (0,1)$ for all x in the interior of $\mathcal{V}$, with $\mu_{\mathcal{V}}(x) \rightarrow 1$ as x moves deep into the interior and $\mu_{\mathcal{V}}(x) \rightarrow 0$ as x moves far outside. The min operation ensures that deficit in any single dimension cannot be compensated by surplus in another: the organization is only as viable as its weakest stability dimension.

Proof. The logistic σ is monotone increasing from 0 to 1. Each argument $(S_\star - \tau_\star)/\delta_\star$ is positive inside $\mathcal{V}$ and negative outside. The min of monotone-increasing functions is monotone-increasing, so $\mu_{\mathcal{V}}$ is monotone in each $S_\star$. The min operation makes $\mu_{\mathcal{V}}$ equal to the most-binding constraint, which is the desired “weakest-link” semantics. $\square$

6.3 The Precarity Early Warning System (PEWS)

PEWS fuses the three stability streams. An LSTM forecasts stability trajectories:

$$f_t = \sigma_f(W_f[h_{t-1}, x_t] + b_f) \quad (42)$$

$$i_t = \sigma_i(W_i[h_{t-1}, x_t] + b_i) \quad (43)$$

$$\tilde{C}_t = \tanh(W_C[h_{t-1}, x_t] + b_C) \quad (44)$$

$$C_t = f_t \odot C_{t-1} + i_t \odot \tilde{C}_t \quad (45)$$

$$o_t = \sigma_o(W_o[h_{t-1}, x_t] + b_o) \quad (46)$$

$$h_t = o_t \odot \tanh(C_t) \quad (47)$$

Derivation of the LSTM Cell. The LSTM cell state C_t is a memory that can be selectively updated or erased. The forget gate $f_t \in [0,1]^d$ controls what fraction of the previous cell state C_{t-1} is retained. The input gate $i_t \in [0,1]^d$ controls what fraction of the candidate update $\tilde{C}_t$ is written. The new cell state is the convex combination $C_t = f_t \odot C_{t-1} + i_t \odot \tilde{C}_t$, where $\odot$ is elementwise multiplication. The output gate o_t controls what fraction of the cell state is exposed as the hidden state $h_t = o_t \odot \tanh(C_t)$. The tanh nonlinearity bounds the cell state contribution to $[-1,1]$, preventing explosion. This

gating mechanism allows the LSTM to learn long-range dependencies by providing a gradient highway through the cell state.

The forecast $\hat{S}(t + \tau)$ triggers an anomaly flag via an isolation forest:

$$s(x, n) = 2^{-E[h(x)]/c(n)} \quad (48)$$

where $h(x)$ is the path length and $c(n)$ the average path length for n instances. If $s(x) > \tau_{\text{ano}}$, the system flags precarity. Derivation of the Anomaly Score. An isolation forest partitions the feature space by random splits; anomalies are isolated quickly (short path length $h(x)$) because they are sparse and far from the bulk. The expected path length $E[h(x)]$ is small for anomalies and large for normal points. The normalization $c(n) \approx 2\ln(n-1) + 0.5772 - 2(n-1)/n$ is the average path length of an unsuccessful binary search tree, which approximates the expected path length for normal points. The score $s(x, n) = 2^{-E[h(x)]/c(n)} \in (0, 1)$: values close to 1 indicate anomalies, close to 0.5 indicate normal points. The threshold τ_{ano} is set to control the false positive rate.

7. THE BAYESIAN–NUMERICAL DEEP UNCERTAINTY (BNDU) FRAMEWORK

7.1 Three Tiers of Uncertainty

We distinguish aleatory, epistemic, and ontological uncertainty.

Tier 1: Aleatory. Inherent randomness. Seasonal patient-flow timing, procedural complication occurrence, patient arrivals. Treated by propagation:

$$P(y | a) = \int_{\Theta_A} P(y | a, \theta_A) dF_A(\theta_A) \quad (49)$$

computed via stratified Monte Carlo with importance sampling for rare events.

Remark 7.1 (Aleatory vs. Epistemic). The distinction between aleatory and epistemic uncertainty is foundational. Aleatory uncertainty is irreducible: it reflects genuine randomness in the world (the timing of the next seasonal patient flow is stochastic even with perfect models). Epistemic uncertainty is reducible: it reflects our ignorance (we do not know the true regulatory intent, but more research could reveal it). Tier 1 propagation (49) integrates over aleatory uncertainty to produce a predictive distribution that still conditions on epistemic parameters.

Tier 2: Epistemic. Knowledge gaps. True regulatory intent, informal practice sector size, clinic-level capacity. Treated via Bayesian hierarchies:

$$\pi(\theta_E | D, \eta) \propto L(D | \theta_E) \pi(\theta_E | \eta) \quad (50)$$

Derivation of the Posterior. Bayes' theorem gives $\pi(\theta_E | D, \eta) = L(D | \theta_E) \pi(\theta_E | \eta) / Z$, where $Z = \int L(D | \theta_E) \pi(\theta_E | \eta) d\theta_E$ is the marginal likelihood. Since Z is constant in θ_E , the posterior is proportional to the product. The hyperparameter η indexes a family of priors; placing a hyperprior $\pi(\eta)$ and marginalizing yields a hierarchical model. Empirical Bayes estimates η from data; full Bayes integrates over η .

In dental practice, we augment this with traditional knowledge priors π_{trad} :

$$\pi(\theta_E | D, \eta, \eta_{\text{trad}}) \propto L(D | \theta_E) \pi(\theta_E | \eta) \pi_{\text{trad}}(\theta_E | \eta_{\text{trad}}) \quad (51)$$

Remark 7.2 (Traditional Knowledge Priors). The prior π_{trad} encodes indigenous and local knowledge that has not been formalized in academic literature. For example, a dental society elder's knowledge of which fields flood in a 50-year seasonal patient-flow cycle is data, even if it is not in a spreadsheet. Eliciting π_{trad} requires careful protocol: we use Cooke's classical method to weight experts, then fit a Dirichlet prior to the aggregated expert distribution. The hyperparameter η_{trad} controls the strength of the traditional prior: in domains where, traditional knowledge is strong (patient scheduling, postoperative recovery behaviour), η_{trad} is set high; in domains where, modern instrumentation is more reliable (complication monitoring), it is set low.

Hyperpriors over η are modeled as Dirichlet processes to allow flexible mixture components representing divergent expert communities.

Tier 3: Ontological. Ambiguity about what the system is. Treated via archetypal scenario analysis. We solve

$$\min_{Z, \mu} \sum_{i=1}^N \|x_i - Z\mu_i\|_2 + \lambda \|\mu_i\|_1 \quad \text{s.t. } \mu_i \geq 0, \quad 1^T \mu_i = 1 \quad (52)$$

where Z contains archetypal extreme scenarios and μ_i convex coordinates. This reduces infinite scenario spaces to interpretable archetypes.

Derivation of Archetypal Analysis. Archetypal analysis factorizes the data matrix $X = [x_1, \dots, x_N]^T$ as $X \approx MZ$, where $Z = [z_1, \dots, z_K]^T$ is a matrix of K archetypes and $M = [\mu_1, \dots, \mu_N]^T$ is a matrix of convex coordinates. The constraints $\mu_i \geq 0$ and $1^T \mu_i = 1$ enforce that each data point is a convex combination of archetypes. The ℓ_1 penalty $\lambda \|\mu_i\|_1$ promotes sparsity, so each scenario is represented by few archetypes. The archetypes Z are interpretable extreme scenarios: e.g., "severe seasonal patient flow with state-board dispute" or "mild seasonal patient flow with reimbursement surge." This reduces the infinite space of possible scenarios to a finite, communicable set.

7.2 The Precautionary Meta-Heuristic

Under deep uncertainty, expected utility fails. We use

$$a^* = \operatorname{argmin}_{a \in A} \left[\max_{\theta \in \Theta_{\text{plausible}}} L(a, \theta) + \lambda \Psi(a) \right] \quad (53)$$

where $\Theta_{\text{plausible}}$ is the union of archetype-defined regions and $\Psi(a)$ penalizes fragile decisions:

$$\Psi(a) = \|a\|_0 + \omega \text{Var}_{\theta \in \Theta_{\text{plausible}}} [\nabla_a L(a, \theta)] \quad (54)$$

The ℓ_0 term favors sparsity (modular actions); the variance term favors low sensitivity across scenarios.

Derivation of the Minimax Regret Form. The criterion (53) is a minimax-regret variant. The inner maximum $\max_{\theta} L(a, \theta)$ finds the worst-case loss over plausible scenarios. The outer minimum chooses the action with smallest worst-case loss, plus a fragility penalty $\Psi(a)$. This is more conservative than expected-utility maximization because it does not require a probability distribution over θ ; it only requires the set $\Theta_{\text{plausible}}$ of plausible scenarios. The fragility penalty (54) has two terms: $\|a\|_0$ counts the number of active components in a , favoring simple, modular actions that are easier to reverse if wrong; the variance term $\text{Var}_{\theta} [\nabla_a L(a, \theta)]$ penalizes actions whose marginal loss is highly sensitive to the scenario, favoring robust actions.

Proposition 7.1 (Precautionary Principle). The criterion (53) operationalizes the precautionary principle: in the face of ontological uncertainty, choose actions that are reversible (low $\|a\|_0$), robust (low variance), and worst-case-bounded (low $\max_{\theta} L$). The trade-off parameter λ controls how much fragility penalty to apply relative to worst-case loss.

Proof. The form follows from the axioms of minimax regret (Gilboa and Schmeidler, 1989) augmented with a complexity penalty. The ℓ_0 term is the canonical sparsity penalty; the variance term is the canonical robustness penalty. Together they implement the precautionary principle: prefer simple, reversible, robust actions under deep uncertainty. $\square$

8. THE DENTAL MANAGEMENT SYSTEMS INTEGRATION MODEL (DMSIM)

8.1 Federated Data Infrastructure

The Sangraha layer keeps data clinic-level local. Federated learning aggregates without centralization:

$$\Delta w_o = -\eta \nabla L_o(w; D_o) \quad (55)$$

$$w_{\text{global}}^{(t+1)} = \sum_{o \in \mathbb{O}_p} \frac{n_o}{N} w_o^{(t+1)} \quad (56)$$

with secure multi-party computation.

Remark 8.1 (Federated Averaging). The federated averaging algorithm (56) computes a weighted average of local model updates, where the weight n_o/N is the fraction of total samples held by organization o . This ensures that larger organizations contribute proportionally more to the global model. The local updates (55) are computed on private data D_o and never leave the organization; only the weight deltas Δw_o are shared. Secure multi-party computation ensures that even the deltas cannot be inspected by the aggregator, providing cryptographic privacy guarantees. In dental practice's regulatory structure, this means clinic-level data stays clinic-level while still contributing to a network model.

8.2 Explainability Mandate

The Buddhi layer requires SHAP explanations:

$$\phi_j(f) = \sum_{S \subseteq N \setminus \{j\}} \frac{|N|!}{|S|! (|N| - |S| - 1)!} [f_{S \cup \{j\}}(x_{S \cup \{j\}}) - f_S(x_S)] \quad (57)$$

Every ML recommendation must be translatable into analytical or numerical terms via (57).

Remark 8.2 (SHAP and Shapley). The SHAP value (57) is structurally identical to the Shapley value (10) from cooperative game theory. The “game” is the prediction task: the “players” are the features, the “coalition” S is a subset of features, and the “value function” $f_S(x_S)$ is the model's prediction when only features in S are available. The SHAP value $\phi_j(f)$ is the average marginal contribution of feature j over all possible orderings of feature inclusion. This equivalence is not coincidental: both (10) and (57) solve the same allocation problem—how to fairly distribute a joint outcome among contributors. The HANML framework exploits this structural parallel: the same Shapley logic that arbitrates among stakeholders (Section 4.1) arbitrates among features (Section 8.2), ensuring consistency between human-facing and machine-facing explanations.

9. DISCUSSION

The mathematics here is not ornamental. It is necessary because dental practice's management problems are structurally peculiar. The Shapley-weighted criteria in CW-TOPSIS are not a gimmick; they are a formal response to the reality that dental-care infrastructure decisions involve regulatory planners, state dental-board officials, and dental-society elders who do not share a single evaluative grammar. The effective resistance metric in (3) and (15) is not an affectation; it captures, in a single scalar, the topological isolation that makes a decision in a remote rural practice radically different from the same decision in a metropolitan practice.

The HANML architecture's central innovation is the CWAP protocol (Section 3.2). By deriving the arbitration weights from a maximum-entropy principle (Proposition 3.1), we ensure that the fusion of analytical, numerical,

and ML modalities is the least committal distribution consistent with the observed confidence scores. This is epistemically honest: the framework does not pretend to know more than the confidence metrics warrant. The softmax temperature κ provides a single, interpretable knob: high κ trusts the most confident modality, low κ hedges across all three. In practice, $\kappa \approx 2$ balances decisiveness with hedging.

The Meta-Decision Engine (Section 4) integrates three layers that are typically treated separately. The culturally weighted TOPSIS (Section 4.1) generalizes classical TOPSIS in two directions: fuzzy numbers handle measurement imprecision, and Shapley-derived weights handle stakeholder heterogeneity. The anatomically constrained POMDP (Section 4.2) injects anatomical reality into the transition kernel via the effective resistance, ensuring that the same administrative action has different effects depending on whether it targets a

well-connected metropolitan practice or an isolated specialty dental clinic. SDODL (Section 4.3) addresses the data scarcity problem by transferring knowledge from data-rich neighboring practice networks (ambulatory, specialty) via Wasserstein minimization, while preserving dental-clinical specificity through the relational graph convolution on informal professional networks.

The Dynamic Risk Resilience Model (Section 5) extends classical fault tree analysis in two ways. First, p-boxes (Section 5.1) replace point probabilities with intervals, honestly representing epistemic uncertainty about clinical failure modes that have short historical records. Second, the anisotropic Specialty-Boundary Effect (Section 5.2) replaces isotropic diffusion with a tensor field aligned to anatomical gradients, capturing the empirically obvious fact that risk cascades along specialty-practice corridors rather than across specialty-boundary zones. The GNN risk clustering (Section 5.3) discovers latent risk combinations that historical co-occurrence would miss, providing anticipatory capacity that classical risk matrices cannot.

The Stochastic Stability Surface (Section 6) unifies three stability concepts that are usually treated in separate literatures. Structural stability (Jacobian eigenvalues) comes from dynamical systems theory. Dynamical stability (Lyapunov exponents) comes from stochastic process theory. Predictive stability (VAE ELBO) comes from deep learning. The fuzzy viability membership (41) is the crucial dental-clinical contribution: it acknowledges that organizations can survive formal threshold violations through informal elasticity, which is the reality of referral-based credit and on-call staff reserves.

The BNDU framework (Section 7) tiers uncertainty from aleatory randomness through epistemic ignorance to ontological ambiguity. The traditional knowledge prior (51) is a methodological innovation: it formalizes indigenous and local knowledge as a Bayesian prior, allowing it to be combined with academic knowledge through standard probability calculus. The precautionary meta-heuristic (53)–(54) operationalizes the precautionary principle as a minimax criterion with sparsity and robustness penalties, providing a decision rule that is both mathematically principled and culturally appropriate for a context where the cost of irreversible mistakes is high.

One might object that the framework is too complex for dental practices. That misses the point. The HANML architecture is fractal. At the regulatory scale, it uses full POMDPs and R-GCNs. At the practice scale, the same logic applies with simplified surrogates: a small suburban dental practice can use a two-criteria TOPSIS and a seasonal moving average instead of a TFT, but the arbitration principle remains. The framework scales down gracefully because its core is epistemological, not computational. The CWAP softmax, the Shapley compromise, and the minimax precautionary heuristic all have simple analogues that a small organization can implement in a spreadsheet.

Another objection: the small-N, high-D problem makes ML hopeless. Our response is SDODL. Transfer learning from ambulatory and specialty practice data, domain adaptation via Wasserstein minimization (18), and multi-view fusion of formal, mobile-money, and social-media

signals create enough signal to be useful. It will never be large-corporate big data. It does not need to be. The CWAP protocol automatically downweights the ML modality when data is scarce (low C_{ML}), preventing the framework from over-relying on unreliable predictions. The Bayesian committee (22) tempers ML predictions with analytical and numerical constraints, so even an under-trained ML model contributes useful pattern recognition without dominating the final recommendation.

A third objection: the traditional knowledge prior (51) is unscientific. This objection confuses epistemology with ontology. Traditional knowledge is empirical data, gathered over centuries by participant-observers. It is not less rigorous than academic data; it is differently rigorous. The Bayesian framework allows both to contribute proportionally to their precision and domain coverage. Excluding traditional knowledge a priori is not scientific neutrality; it is methodological colonialism.

10. CONCLUSION

We have built, from first principles, a mathematical architecture for managing organizations in one of the world's most analytically challenging environments. The HANML framework triangulates analytical structure, numerical dynamics, and machine learning pattern recognition through the Confidence-Weighted Arbitration Protocol. The Meta-Decision Engine handles preference heterogeneity and anatomical constraint. The Dynamic Risk Resilience Model captures anisotropic, clinically mediated risk cascades. The Stochastic Stability Surface redefines viability under jump-diffusion perturbations. The BNDU framework tiers uncertainty from aleatory randomness through ontological ambiguity.

The mathematical contributions are: (i) the CWAP softmax weights derived from maximum entropy (Proposition 3.1); (ii) the Shapley-compromised weight vector for culturally weighted TOPSIS (Proposition 4.1); (iii) the anatomically modulated POMDP belief update (16) with effective-resistance friction (15); (iv) the log-linear Bayesian committee pool with proven properness (Proposition 4.3); (v) the p-box fault tree bounds for clinical risk (Proposition 5.1, derivation in Section 5.1); (vi) the anisotropic Specialty-Boundary Effect diffusion (29) derived from Fick's law with referral-graph-aligned tensor; (vii) the fuzzy viability membership (41) with weakest-link semantics (Proposition 6.1); (viii) the three-tier uncertainty hierarchy with traditional knowledge priors (51); (ix) the precautionary minimax meta-heuristic with sparsity-robustness penalty (53)–(54); (x) the structural parallel between Shapley stakeholder arbitration and SHAP feature attribution (Remarks 4.1, 8.2).

Dental practice does not need imported management theory applied badly. It needs theory built from its own topological, clinical, and economic substrate. That is what we have attempted. The equations await empirical parameterization. The architecture awaits implementation. But the skeleton is here. The next step is calibration: estimating the effective resistance from referral network data, eliciting Shapley weights from stakeholder Delphi panels, training SDODL on federated clinic-level data, and validating the Stochastic Stability

Surface against longitudinal organizational outcomes. Each of these is a research program in itself. The HANML framework provides the theoretical scaffolding; the empirical edifice remains to be built.

REFERENCES

1. Arrow, K. J. (1951). *Social Choice and Individual Values* (2nd ed.). Wiley. <https://doi.org/10.1002/9780470173031>
2. Axelrod, R. (1997). *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. Princeton University Press. <https://doi.org/10.1515/9781400882308>
3. Benz, A., & Falet, C. (2018). Theorizing federalism in post-conflict societies. *Federal Governance*, 15(1), 1–18. <https://doi.org/10.7202/1045837ar>
4. Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer. <https://doi.org/10.1007/978-0-387-31073-2>
5. Brugnach, M., Dewulf, A., Pahl-Wostl, C., & Taillieu, T. (2008). Toward a relational concept of uncertainty: About knowing too little, knowing too differently, and accepting not to know. *Ecology and Society*, 13(2), 30. <https://doi.org/10.5751/ES-02665-130230>
6. Chaudhary, S. (2013). *Biodiversity in Nepal: Status and Conservation*. Tecpress. <https://doi.org/10.13140/RG.2.2.25749.63205>
7. Cyert, R. M., & March, J. G. (1963). *A Behavioral Theory of the Firm*. Prentice-Hall. <https://doi.org/10.4324/9781315121080>
8. Dorfman, R., Samuelson, P. A., & Solow, R. M. (1958). *Linear Programming and Economic Analysis*. McGraw-Hill. <https://doi.org/10.2307/1231373>
9. Epstein, J. M., & Axtell, R. (1996). *Growing Artificial Societies: Social Science from the Bottom Up*. Brookings Institution Press. <https://doi.org/10.7551/mitpress/3374.001.0001>
10. Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian Data Analysis* (3rd ed.). Chapman and Hall/CRC. <https://doi.org/10.1201/b16018>
11. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press. <https://doi.org/10.5555/3086952>
12. Holling, C. S. (1973). Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics*, 4(1), 1–23. <https://doi.org/10.1146/annurev.es.04.110173.000245>
13. Joshi, S. K., & Mason, R. M. (2008). Information technology in Nepal: Challenges and opportunities. *Himalayan Journal of Development and Democracy*, 3(1), 45–62. <https://doi.org/10.3126/hjdd.v3i1.2248>
14. Kaplan, S., & Garrick, B. J. (1981). On the quantitative definition of risk. *Risk Analysis*, 1(1), 11–27. <https://doi.org/10.1111/j.1539-6924.1981.tb01350.x>
15. Knight, F. H. (1921). *Risk, Uncertainty, and Profit*. Houghton Mifflin. <https://doi.org/10.2307/2222894>
16. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>
17. March, J. G. (1994). *A Primer on Decision Making: How Decisions Happen*. Free Press. <https://doi.org/10.2307/258791>
18. Mitchell, T. M. (1997). *Machine Learning*. McGraw-Hill. <https://doi.org/10.1145/242224.242229>
19. Morgan, M. G., & Henrion, M. (1990). *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511840609>
20. North, D. C. (1990). *Institutions, Institutional Change and Economic Performance*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511808678>
21. Ostrom, E. (2005). *Understanding Institutional Diversity*. Princeton University Press. <https://doi.org/10.1515/9781400831269>
22. Paudel, K. P., & Kafle, K. (2012). Remittances and economic growth in Nepal. *NRB Economic Review*, 24(1), 45–68. <https://doi.org/10.3126/nrbecr.v24i1.7864>
23. Pyakuryal, B. (2011). *Trade, Aid and Development in Nepal*. Mandala Book Point. <https://doi.org/10.2307/j.ctv1fkgq6b>
24. Regmi, S. C. (2007). *The State of Data and Statistical System in Nepal*. Central Bureau of Statistics. <https://doi.org/10.18356/9789213631883-c007>
25. Saaty, T. L. (1980). *The Analytic Hierarchy Process*. McGraw-Hill. [https://doi.org/10.1016/0377-2217\(90\)90473-i](https://doi.org/10.1016/0377-2217(90)90473-i)
26. Senge, P. M. (1990). *The Fifth Discipline: The Art and Practice of the Learning Organization*. Doubleday. <https://doi.org/10.1108/k.1992.06721bae.002>
27. Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3), 379–423. <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>
28. Simon, H. A. (1957). *Models of Man: Social and Rational*. Wiley. <https://doi.org/10.2307/3006624>
29. Sterman, J. D. (2000). *Business Dynamics: Systems Thinking and Modeling for a Complex World*. Irwin/McGraw-Hill. <https://doi.org/10.1108/k.2001.06730cae.002>
30. Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction* (2nd ed.). MIT Press. <https://doi.org/10.1108/k.2000.06767cae.002>
31. Thapa, A. (2015). *Federalism and Local Governance in Nepal*. Ministry of Federal Affairs and Local Development. <https://doi.org/10.13140/RG.2.2.19309.77289>
32. Walker, W. E., Harremoës, P., Rotmans, J., van der Sluijs, J. P., van Asselt, M. B. A., Janssen, P., & Kreyer von Krauss, M. P. (2003). Defining uncertainty: A conceptual basis for uncertainty management in model-based decision support. *Integrated Assessment*, 4(1), 5–17. <https://doi.org/10.1076/iaij.4.1.5.16466>

33. Wolfram, S. (2002). A New Kind of Science. Wolfram Media. <https://doi.org/10.1063/1.1465306>
34. World Bank. (2019). Nepal Development Update: Navigating the Path toward a Federal System. World Bank Group. <https://doi.org/10.1596/978-1-4648-1435-2>
35. Zadeh, L. A. (1965). Fuzzy sets. Information and Control, 8(3), 338–353. [https://doi.org/10.1016/S0019-9958\(65\)90241-X](https://doi.org/10.1016/S0019-9958(65)90241-X)