
Framework Library

29 frameworks, each with a visual model, worked example, usage criteria, limitations and a reusable template.

Product Management Master Handbook
v0.1 · 23 October 2026

Built from the 55-lecture course study guide and extended to the 12-module Product Core curriculum.
Hypothetical examples are labelled; no real company results are invented.

PART

Framework Library

29 frameworks, each with a visual model, worked example, usage criteria, limitations and a reusable template.

30 entries

Framework selection guide: which framework should I use?

Frameworks are thinking tools: they organise evidence and make reasoning visible. They never replace evidence or judgement. Start from **the decision you need to make**, then pick the lightest framework that helps.

If your question is...	Reach for	Why	Avoid
Who is the customer and what are they trying to get done?	JTBD, Customer Journey Map	Frames demand around progress in a situation	Demographic personas alone
Which customer problem should we pursue for our outcome?	Opportunity Solution Tree	Links outcome → opportunities → solutions → tests	Jumping to features
Is our idea worth pursuing at all?	Lean Startup / Build-Measure-Learn, assumption testing	Cheapest test of riskiest assumption	Building the full product first
What should we build first?	RICE / ICE (comparable items); WSJF (time-critical)	Explicit, comparable scoring	Scores as facts; ignoring dependencies
What fits in this release?	MoSCoW	Scope negotiation against a date	Ranking within a category
Which features delight vs are expected?	Kano	Customer expectation categories	Small or unsegmented samples
Choose among few strategic bets	Weighted scoring, Cost of Delay	Encodes strategic criteria and urgency	Reverse-engineered weights
Where are we in the market? Should we enter?	SWOT, PESTLE, Porter's Five Forces	Structured environment scan	Lists without so-what
How does our business create and capture value?	Business Model Canvas, Value Proposition Canvas	Makes the model and fit explicit	Treating cells as facts
What single metric aligns the team?	North Star Framework (+ metric tree), OKRs	Focus with drivers and guardrails	One metric without guardrails
Where is the growth funnel leaking?	AARRR	Stage-by-stage diagnosis	Generic definitions

If your question is...	Reach for	Why	Avoid
How do we measure UX quality?	HEART	Goals → signals → metrics	Tracking all five by default
Did the change work?	Hypothesis testing, A/B testing	Causal evidence with uncertainty	Underpowered tests; peeking
Why did this happen?	Five Whys, fishbone, Pareto	Root causes and focus	Single-chain blame
Are we stuck in assumptions or loops?	First principles, systems thinking	Rebuild from fundamentals; see feedback loops	Over-theorising simple issues
How should we explore a fuzzy problem?	Design Thinking, Double Diamond	Diverge then converge with users	Workshops without real users

Given this problem, which framework? Start from the decision. If a framework doesn't change what you would do, skip it.

Decision tree in words: *Unclear problem?* → discovery frameworks (JTBD, journey, OST, Design Thinking). *Clear problem, many options?* → prioritisation (RICE/ICE/WSJF/MoSCoW/Kano). *Strategic direction?* → SWOT/PESTLE/Porter/canvases. *Measuring?* → North Star, AARRR, HEART, OKRs. *Proving causality?* → hypothesis and A/B testing. *Something broke?* → Five Whys, fishbone, Pareto.

Interview tip

Name a framework only after you have clarified the problem, then *use* it (show a few real scores or boxes) and state its limitation. Interviewers reward applied judgement, not vocabulary.

SWOT

STRENGTHS (INTERNAL, HELPFUL)

- Capabilities, assets, advantages

WEAKNESSES (INTERNAL, HARMFUL)

- Gaps, constraints, liabilities

OPPORTUNITIES (EXTERNAL, HELPFUL)

- Trends, unmet needs, partners

THREATS (EXTERNAL, HARMFUL)

- Competitors, regulation, cost shifts

SWOT model. Follow with TOWS: S×O (use strengths to seize), W×O (fix to seize), S×T (use strengths to defend), W×T (avoid).

Use when: framing a strategy discussion, market entry, annual planning. **Don't use when:** you need prioritised actions or quantified trade-offs — SWOT alone produces lists.

Worked example (hypothetical): regional food-delivery entrant — S: restaurant relationships; W: small marketing budget; O: office-lunch subscriptions; T: incumbents' promo wars. **TOWS action:** S×O → launch office-lunch subscriptions through existing restaurant partners; W×T → don't compete on discounts.

Limitations: subjective, static, often evidence-free; items mix levels of importance. Always attach evidence and rank items.

Template: for each quadrant list ≤ 5 items, each with *evidence* and *importance* (H/M/L); then write 3 TOWS actions with owners. **Deep dive:** Module 2.

PESTLE

POLITICAL

- Government policy, stability, trade

ECONOMIC

- Rates, inflation, income, demand

SOCIAL

- Demographics, behaviour, trust

TECHNOLOGICAL

- Infrastructure, platforms, AI

LEGAL

- Data protection, consumer law, licensing

ENVIRONMENTAL

- Climate, sustainability expectations


PESTLE model. Mark each factor as tailwind or headwind, rate impact × certainty, and assign a watcher.

Use when: entering a new market or country; long-range planning; regulated products. **Don't use when:** you need customer-level insight (use research) or competitor detail (use competitor analysis).

Worked example (hypothetical): fintech lending in a new market — *Legal:* data-protection consent and lending disclosures (high impact, high certainty → design consent flows early); *Technological:* digital ID infrastructure (tailwind → faster KYC); *Economic:* rising interest rates (headwind → stricter credit policy).

Limitations: can become a long list with no prioritisation; factors interact; needs expert input for legal/regulatory items. **Template:** factor · description · tailwind/headwind · impact (1–5) · certainty (1–5) · implication · owner. **Deep dive:** Module 2.

Porter's Five Forces




Five Forces model. Stronger forces mean lower structural profit potential. Ask which force your product can weaken.

Use when: deciding whether to enter a market, explaining margin pressure, positioning against structural threats. **Don't use when:** analysing a single feature, or in fast-changing ecosystems without considering complements and platforms.

Worked example (hypothetical): food delivery — rivalry high (promotions), buyer power high (low switching costs), supplier power medium–high (restaurants multi-home), substitutes high (cooking, dining out), entrants medium (density is hard). **Implication:** invest in habit/loyalty (reduce buyer power) and partner tools (reduce supplier churn).

Limitations: static snapshot; industry boundaries are fuzzy; ignores network effects and complementors. **Template:** force · rating (L/M/H) · evidence · trend · product lever. **Deep dive:**


Business Model Canvas



Use when: describing or redesigning a business model; comparing models; onboarding stakeholders to the economics. **Don't use when:** you're still searching for a problem (use Lean Canvas) or need financial detail (build a model).

Worked example: see the meal-planning app canvas in Module 2. **Limitations:** static; doesn't show competition or sequencing; cells are assumptions until tested. **Template:** the nine headings with *assumption* / *evidence* / *test* for each cell. **Deep dive:** Module 2.

Value Proposition Canvas




Use when: shaping a new offering or repositioning; aligning product and marketing on what matters. **Don't use when:** you have no customer research yet — the right side would be guesses.

Worked example (hypothetical): busy parents — job: feed the family on budget; top pain: surprise grocery bills; reliever: price estimate before shopping; gain creator: weekly budget tracker. **Test:** smoke test with the promise “Know your weekly grocery cost in 5 minutes”.

Limitations: easy to fill with assumptions; doesn't capture competition. **Template:** rank top 3 jobs, pains, gains with evidence; map each reliever/creator to a ranked pain/gain. **Deep dive:** Module 2.

North Star Framework




Use when: aligning several teams on a shared outcome; moving from output to outcome culture. **Don't use when:** a product has distinct, unrelated value streams (use separate goals) or the metric can't be measured reliably yet.

Worked example (hypothetical): food app NSM = completed problem-free orders per active customer per month; inputs = activation, order frequency, payment success, on-time delivery; guardrails = refund rate, complaint rate.

Limitations: can be gamed (Goodhart's law); one metric never tells the whole story; may need revisiting as strategy changes. **Template:** NSM definition (numerator, denominator, window) · why it reflects value · 3–5 input metrics with owners · 2–3 guardrails · review cadence. **Deep dive:** Module 7.

OKRs (Objectives and Key Results)



OKR structure. Shipping an initiative is not achieving a key result; the KR measures the outcome.

Use when: quarterly focus and alignment; outcome-based planning. **Don't use when:** as a task list, an individual performance score, or with 10+ objectives.

Worked example: O: make checkout trustworthy for first-time buyers. KR1: completion 70% → 76% (new users). KR2: payment-failure recovery 20% → 35%. KR3 (guardrail): complaint rate ≤ 1.3%. Initiatives: early fee display; payment retry.

Limitations: sandbagging or unrealistic stretch; metric gaming; overhead. **Template:** Objective · KRs (metric, baseline, target, date, owner) · initiatives · weekly confidence score (0–10). **Deep dive:** Module 2.

Jobs to Be Done (JTBD)

JOB STATEMENT

- When [situation], I want to [motivation], so I can [expected outcome].

PUSH (AWAY FROM TODAY)

- Frustrations with the current way

PULL (TOWARD NEW)

- Attraction of the new solution

ANXIETY (ABOUT NEW)

- Fears and uncertainties

HABIT (OF TODAY)

- Inertia, sunk costs, routines

JTBD model — job statement and the four forces. Switching happens when push + pull > anxiety + habit.

Use when: discovering demand; understanding competition broadly (including non-consumption); positioning. **Don't use when:** you need usability detail (use testing) or market sizing.

Worked example (hypothetical): *When I'm in back-to-back meetings with 15 minutes to eat, I want a filling lunch at my desk without surprise costs, so I can stay focused and on budget.*
Competitors: snacks, colleagues, skipping lunch. Implication: reliability and price clarity beat variety.

How to run: interview recent switchers about the *timeline* of their decision (first thought → passive looking → active looking → decision → first use); map forces; write 2–4 job statements; validate frequency with data.

Limitations: interpretive; job statements can become vague slogans; needs good interviewing.

Template: job statement · situation triggers · desired outcomes (measurable) · current alternatives · push/pull/anxiety/habit evidence. **Deep dive:** Module 3.

Customer Journey Mapping



Journey emotion curve (hypothetical). Rows of a full map: stages · actions · thoughts · emotions · touchpoints · pain points · opportunities. Low points are where to look first.

Use when: aligning teams on the end-to-end experience; finding cross-channel gaps (app + support + delivery). **Don't use when:** you need screen-level flow detail (use a user flow) or have no research to base it on.

Worked example (hypothetical): weekday dinner journey — lowest points at “cart total” (fee surprise) and “support” (slow refunds). Opportunities: show total early; proactive refund status.

Limitations: can become a beautiful poster nobody uses; averages hide segment differences.

Template: columns = stages; rows = actions, thoughts/quotes, emotion score, touchpoints, pains, opportunities, owners, metrics. **Deep dive:** Modules 3 and 5.

Opportunity Solution Tree




OST model. Compare opportunities before solutions; test multiple solutions per opportunity.

Use when: continuous discovery with a clear outcome; avoiding solution-first thinking; communicating why you're working on something. **Don't use when:** you lack an agreed outcome, or the work is mandatory (compliance).

Worked example: outcome = checkout completion 70% → 76%; opportunities: fee surprise, payment-failure dead ends, long first-order form; solutions and tests as in Module 3.

Limitations: requires ongoing research to stay true; can grow unwieldy. **Template:** outcome · opportunities (from interviews, with evidence) · ≥ 3 solutions per chosen opportunity · assumption tests. **Deep dive:** Module 3.

Design Thinking



Design Thinking stages. Non-linear; testing often loops back to Define or Ideate.

Use when: ambiguous problems, new experiences, cross-functional workshops. **Don't use when:** the problem and solution are well understood and the risk is delivery, not discovery.

Worked example: checkout redesign — empathise (interviews: fee surprise) → define ("first-time buyers can't see total before paying") → ideate (8 concepts) → prototype (2 Figma variants) → test (6 users; task success 50% → 83%).

Limitations: workshops without real users become theatre; can neglect business viability and feasibility. **Template:** problem statement · research insights · how-might-we questions · concepts · prototype plan · test tasks · learnings. **Deep dive:** Module 5.

Double Diamond



Use when: planning a discovery-to-delivery project and explaining to stakeholders why time goes into problem definition. **Don't use when:** for small iterative changes.

Worked example: a fintech onboarding project spends 3 weeks in Discover (interviews, funnel analysis), converges on "document upload fails on low-end phones," then explores 5 solutions and delivers camera guidance + save-and-resume.

Limitations: looks linear; real work loops. **Template:** for each diamond: activities, outputs, decision gate, owner. **Deep dive:** Module 5.

Lean Startup

VALIDATED LEARNING

- Progress = what you learned with evidence

MVP

- Smallest experiment to test a key assumption

INNOVATION ACCOUNTING

- Baseline → tune → measure progress on actionable metrics

PIVOT OR PERSEVERE

- Change strategy based on evidence, not ego

ACTIONABLE VS VANITY METRICS

- Cohort-based, causal metrics over totals

SMALL BATCHES

- Ship small, learn fast

Lean Startup principles. Popularised by Eric Ries (The Lean Startup, 2011).

Use when: high uncertainty about customer, problem, solution or business model. **Don't use**


when: in highly regulated or safety-critical contexts where an "MVP" cannot be released without full assurance (use prototypes and simulations instead).

Worked example: meal-planning app — leap-of-faith assumptions: parents plan weekly; trust price estimates; will pay ₹1,200/yr. MVP sequence: fake door (demand) → concierge (value) → priced pilot (willingness to pay). Pivot if pilot conversion < 3%.

Limitations: misread as "ship low quality"; startup framing may not fit enterprise constraints.

Template: assumptions ranked by risk · experiment per assumption · success threshold · decision. **Deep dive:** Module 4.

Build–Measure–Learn




Build–Measure–Learn loop. Plan in reverse: decide what you need to learn, what to measure, then what to build.

Use when: iterating toward fit; structuring sprints around learning goals. **Don't use when:** measurement is impossible within a useful time frame (use qualitative tests instead).

Worked example: learn whether early fee display reduces surprise → measure task success in 6 usability sessions and later checkout completion in an A/B test → build a two-screen prototype, then a flagged feature.

Limitations: “Build” first framing can encourage building before learning; loops need clear decision rules. **Template:** learning question · metric and threshold · smallest build · result · decision (pivot/persevere/stop). **Deep dive:** Modules 4 and 7.

RICE



RICE scores for a hypothetical backlog. From datasets/rice_backlog.csv (hypothetical). Run a sensitivity check before trusting the order.

RICE = Reach × Impact × Confidence ÷ Effort

Impact scale: 3 massive · 2 high · 1 medium · 0.5 low · 0.25 minimal


Confidence: 100% high · 80% medium · 50% low Effort: person-months (> 0)

Use when: many comparable items, some data on reach. **Don't use when:** comparing mandatory work, platform investments with diffuse reach, or items with unknowable impact.

Worked example: fee early = $40,000 \times 2 \times 0.8 \div 2 = 32,000$; voice ordering = $8,000 \times 2 \times 0.3 \div 8 = 600$ (see Module 4 for all six and the stakeholder script).

Limitations: false precision; bias toward high-reach items; confidence is subjective. **Template:** feature · reach (unit, period, source) · impact (scale + rationale) · confidence (evidence) · effort (eng estimate) · score · notes on dependencies/risks. **Deep dive:** Module 4.

ICE



ICE scores (hypothetical). $I \times C \times E$ on 1–10 scales. Some teams average the three instead of multiplying.

Use when: quick triage of experiment ideas with similar audiences. **Don't use when:** reach varies widely (use RICE) or decisions are high-stakes and need more rigour.

Worked example: fee early $I8 \times C8 \times E7 = 448$; voice $I7 \times C3 \times E2 = 42$.

Limitations: very subjective; easy to game; ignores reach and dependencies. **Template:** idea · hypothesis · I (why) · C (evidence) · E (effort note) · score · owner. **Deep dive:** Module 4.

MoSCoW

MUST

- Release fails without it

SHOULD

- Important, painful to omit

COULD

- Nice if time allows

WON'T (THIS TIME)

- Explicitly out of scope


MoSCoW model. Agree the release period and why each item sits where it does.

Use when: fixed-date releases, contract scope, MVP definition. **Don't use when:** ranking many items against each other (it doesn't order within categories).

Worked example: early-fee release — Must: estimate + labels + fallback; Should: refresh on address change; Could: “add ₹X to avoid fee” nudge; Won't: fee-policy change.

Limitations: “Must” inflation; no numeric basis. Guideline: keep Musts to roughly 60% or less of capacity. **Template:** item · category · rationale · owner · dependency. **Deep dive:** Module 4.

Kano Model




Use when: deciding what makes a competitive baseline vs a differentiator; planning releases that balance basics and delight. **Don't use when:** you lack enough respondents per segment, or need cost/feasibility trade-offs.

Worked example: payment confirmation = must-be; faster checkout = performance; surprise free dessert = attractive; dark mode = indifferent for most, performance for night-shift users.

Limitations: survey wording sensitivity; categories shift; ignores cost. **Template:** feature · functional question · dysfunctional question · answer distribution · category (see the evaluation table in Module 4). **Deep dive:** Module 4.

Weighted Prioritization (Scoring)



Weighted scores (hypothetical). Weights: impact 35%, revenue 25%, strategic fit 20%, ease 20%.


$$\text{Score} = \sum (\text{weight}_i \times \text{score}_i) \quad \text{weights sum to 100\%}$$

Use when: choosing among a few strategic initiatives where criteria differ from pure reach/effort. **Don't use when:** weights are set after seeing scores, or with too many criteria.

Worked example: payment retry $5 \times .35 + 4 \times .25 + 4 \times .20 + 3 \times .20 = 4.15$ (Module 4 has all four).

Limitations: subjective scores; hides trade-offs inside a total. Test sensitivity by shifting weights ± 10 points. **Template:** criteria + weights (agreed first) · anchored 1–5 scales · scores with rationale · total · sensitivity notes. **Deep dive:** Module 4.

Cost of Delay



Cumulative value lost by delaying a feature worth ₹50,000/week (hypothetical). Linear urgency profile. Others include fixed-date (value drops to zero after a deadline) and decaying (first-mover advantage).

Cost of Delay = value per unit time × delay


CD3 = Cost of Delay ÷ duration (do highest first)

Use when: arguing for capacity, sequencing time-sensitive work, comparing urgent vs important. **Don't use when:** value per time can't be estimated even roughly.

Worked example: a 6-week delay of a ₹50,000/week feature costs **₹3,00,000** — compared with the ₹1,20,000 cost of a contractor to accelerate it.

Limitations: estimates are uncertain; urgency profiles differ. **Template:** item · value/week (range) · urgency profile · expected delay · CoD · duration · CD3. **Deep dive:** Module 4.

WSJF (Weighted Shortest Job First)



WSJF ranking (hypothetical). Relative scales (e.g., 1, 2, 3, 5, 8, 13, 20). Scale each component across all jobs at once.


$$\text{WSJF} = (\text{User/business value} + \text{Time criticality} + \text{Risk reduction/opportunity enablement}) \div \text{Job size}$$

Use when: scaled-agile planning with many teams and competing epics. **Don't use when:** job sizes are unknown or political.

Worked example: C: $(5 + 13 + 2) \div 3 = 6.67$; B: $(13 + 8 + 5) \div 13 = 2.0$ — the small urgent fix goes first.

Limitations: relative scores are subjective; can starve big strategic work. **Template:** job · value · time criticality · RR/OE · CoD · size · WSJF. **Deep dive:** Module 4.

AARRR (Pirate Metrics)



AARRR as a funnel (hypothetical monthly numbers). Hypothetical. Stages can overlap in time; referral can feed acquisition in a loop.

Use when: diagnosing a growth problem; organising a metrics dashboard; growth-team planning. **Don't use when:** applied with generic definitions — each stage must be defined for *your* product.

Worked example: activation (42%) and 30-day retention (45% of activated) are the biggest levers; improving activation by 5 points adds ~500 activated users/month at the same acquisition spend.

Limitations: linear view of a non-linear journey; ignores user value quality; no guardrails by default. **Template:** stage · event definition · metric · current · target · owner · top hypothesis.

Deep dive: Module 7.

HEART

Dimension	Goal	Signal	Metric
Happiness	Users find search satisfying	Post-search rating	Avg rating; % positive
Engagement	Users rely on search	Searches per active user	Weekly searches per user
Adoption	New users try search	First search	% new users searching in week 1
Retention	Users keep searching	Repeat search	% searchers searching again next week
Task success	Users find what they want	Click/add-to-cart after search	Search success rate; time to result

HEART with Goals–Signals–Metrics (example: search feature). Pick the dimensions that match your goal; you rarely need all five.


Use when: defining UX success for a feature or redesign. **Don't use when:** you need business-outcome metrics only (pair with North Star/AARRR).

Worked example: checkout redesign — Task success (completion rate, errors), Happiness (CSAT), Adoption (first-time buyer completion).

Limitations: survey-based happiness is noisy; engagement can conflict with wellbeing.

Template: dimension · goal · signal · metric · baseline · target. **Deep dive:** Module 5.

Hypothesis Testing



Two-sided test at $\alpha = 0.05$. The course example (23/100 vs 37/100) gives $z \approx 2.16$, $p \approx 0.031$ — significant at 0.05, not at 0.01.


$$z = (p_2 - p_1) \div \sqrt{\hat{p}(1 - \hat{p})(1/n_1 + 1/n_2)}, \quad \hat{p} = (x_1 + x_2)/(n_1 + n_2)$$
$$p\text{-value} = P(\text{result at least this extreme} \mid H_0 \text{ true})$$

Use when: deciding whether an observed difference is likely to be more than noise. **Don't use when:** samples are tiny, assignment isn't random, or you're fishing across many metrics without correction.

Worked example: 7,000/10,000 vs 7,250/10,000 $\rightarrow z = 3.91$, $p \approx 0.00009$, 95% CI +1.25 to +3.75 pts (Module 7).

Limitations: p-values are widely misread (not the probability the variant works); significance $\neq$ importance. **Template:** $H_0/H_1 \cdot \alpha \cdot \text{test} \cdot \text{sample size/power} \cdot \text{result (effect, CI, p)} \cdot \text{decision}$. **Deep dive:** Module 7.

A/B Testing




Use when: enough traffic, reversible change, measurable short-term outcome. **Don't use when:** low traffic, strong network effects/contamination, legal or safety changes, or long-horizon outcomes.

Worked example: sizing for 70% → 72% at α 0.05 and 80% power: $\approx$ **8,077 users per arm** (Module 7).

Limitations: novelty effects, interference between users, multiple-testing false positives, only measures what you instrument. **Template:** hypothesis · population · variants · primary metric · guardrails · MDE · n per arm · duration · analysis plan · decision rule. **Deep dive:** Module 7.

Five Whys




Use when: post-incident reviews, recurring defects, metric drops. **Don't use when:** problems have many interacting causes (use fishbone/systems thinking) or the exercise turns into blame.

Worked example: late orders → riders wait → food not ready → optimistic prep estimates → estimates ignore kitchen load → add live load signal (Module 4).

Limitations: single causal chain; depends on knowledge in the room. **Template:** problem · why 1–5 with evidence · root cause · countermeasure · owner · verification. **Deep dive:** Module 4.

Pareto Analysis



Complaint categories (synthetic, 100 complaints). The top four categories account for 87% of complaints.

Use when: many issues compete for attention; support or defect triage. **Don't use when:** severity varies hugely (weight by impact, not count) or categories are poorly defined.

Worked example: fixing late delivery and wrong items addresses 60% of complaints; but payment failures may be more severe per case — weight by cost or harm before deciding.

Limitations: "80/20" is a heuristic, not a law; frequency $\neq$ importance. **Template:** category · count · % · cumulative % · severity weight · owner. **Deep dive:** Module 4.

First Principles Thinking




First-principles process. Useful when an industry "standard" may be habit rather than necessity.

Use when: stuck in incremental thinking; cost structures seem fixed; designing something new.
Don't use when: the problem is routine and well-solved — conventions (Jakob's law) exist for good reasons.

Worked example: "30-minute delivery promise because competitors promise it" → fundamentals: customers want reliability; costs are travel and wait → promise an honest live-load window instead.

Limitations: time-consuming; risk of ignoring hard-won conventions. **Template:** goal · assumption list · evidence for each · which are real constraints · new design options. **Deep dive:** Module 4.

Systems Thinking



Balancing loop — notifications and fatigue (conceptual). The more you push, the more fatigue limits future reach. Guardrail: opt-out rate.

Use when: metrics move in unexpected directions; marketplaces; growth tactics with side effects; policy changes. **Don't use when:** a simple, isolated bug (use Five Whys).

Worked example: surge pricing increases supply (reinforcing for availability) but raises cancellations and mistrust (balancing) — add caps and fare transparency.

Limitations: diagrams can become complex; hard to quantify. **Template:** variables · links (+/–) · loops (R/B) · delays · leverage points · guardrail metrics. **Deep dive:** Module 4.