

Travel Analytics portfolio

Four evidence-governed cases for content quality, funnel decisions and operating workflow

Independent portfolio work. Public-source, synthetic, proxy and simulated evidence stay visibly labelled throughout.

PUBLIC / DERIVED

SYNTHETIC / SIMULATED

PROXY / UNAVAILABLE

Evidence map

Each case answers a different decision. The evidence types are not interchangeable.

PROJECT 1	Public source + derived + proxy	Content triage
PROJECT 2	Synthetic clickstream + simulated experiment	Ship decision
PROJECT 3	Mixed upstream package + simulated workflow	Weekly operating view
PROJECT 4	Public source + derived + simulated workflow	Content operations

Project 1: content quality triage

PUBLIC SOURCE + DERIVED

Which listings need content remediation first?

23,233 Bangkok listings scored after cleaning. 11,480 listings, 49.41%, fall below the provisional score threshold of 70.

Decision use
Review highest-priority content gaps
through the RICE backlog.

Project 1: proxy boundary

MODELLED REVIEW PROXY

8,908 estimated additional reviews per year for the top 50 candidates

The estimate helps sequence a content-review queue. It does not establish bookings, revenue or causal impact from content changes.

What would be needed

Observed outcome data and a credible experiment or quasi-experimental design.

Project 2: synthetic funnel

SYNTHETIC CLICKSTREAM

Should a review-prominence PDP badge ship?

25,000 synthetic sessions · 77,096 synthetic events · 722 purchases · 2.89% overall CVR

Funnel diagnosis points to the view_listing to select_dates drop. Project 1 association motivates a hypothesis only; it does not prove demand-side impact.

Project 2: simulated experiment

SIMULATED PRIMARY READOUT

DO_NOT_SHIP

Control 3.64% · treatment 3.66% · relative uplift +0.69% · $p=0.936120$ · 95% CI $[-0.59\text{pp}, +0.64\text{pp}]$

The simulation is underpowered against the design. Exploratory device segments cannot overrule the primary decision.

Project 3: weekly operating package

MIXED UPSTREAM EVIDENCE

How can content and funnel evidence become a weekly-style operating package?

Project 3 reconciles Project 1 content metrics and Project 2 funnel and experiment artifacts while retaining each upstream truth label.

Action

Keep content triage moving and retain the simulated DO_NOT_SHIP decision.

Project 3: operating limits

ILLUSTRATIVE WORKFLOW

Weekly package limits

Illustrative WoW compares demonstration snapshots only. It is not a live operating trend. Placeholder owners and channels are not workplace teams.

Unavailable claim

No reproducible timing evidence supports the manual-process time-saving claim, so it is omitted from public materials.

Project 4: public-source content operations

PUBLIC SOURCE + DERIVED

How should public-source content issues enter a review workflow?

31,069 Inside Airbnb Bangkok listings, snapshot 2026-06-29. 14,551 listings, 46.8%, sit below provisional readiness 70.

Action

Run a two-lane remediation pilot and keep entity matches in human review.

Project 4: simulated workflow limits

DETERMINISTIC SIMULATED WORKFLOW

13,178 open backlog · 896 critical open · 21.73% SLA

These figures illustrate ticket/SLA workflow fields. They do not measure a real SLA, team performance or production operations.

Priority uses review activity as a demand proxy. Entity matches stay review candidates until a labelled sample supports precision/recall evaluation.

How the four systems connect

Project 1 identifies content-risk candidates. Project 2 tests a demand-side hypothesis with synthetic data. Project 3 packages the evidence into a weekly workflow. Project 4 demonstrates public-source content operations and review controls.

Shared rule

Claims keep their original evidence class as they move between artifacts.

Validation and remaining limits

LOCAL INTEGRATION COMPLETE

What the portfolio demonstrates

Inspectable data contracts, reproducible artifact checks, responsive dashboards, and recruiter-facing disclosure discipline.

What remains unproven

Agoda production outcomes, bookings, revenue, causal impact, workplace SLA performance, P3 timing savings, and P4 entity-match precision/recall.