From Strategic Foresight to Inspectable Decisions

A Provenance-Aware Scenario Console for Thailand's Economic Clusters

Poon Thiengburanathum and Non Arkaraprasertkul

Research: Poon Thiengburanathum. **Analysis:** Non Arkaraprasertkul.

Preprint draft, version 1, 5 November 2026. Not peer reviewed.

Abstract

Strategic foresight combines observations, judgements and conditional recommendations, but an interactive presentation can make these different claims look equally measured. We examine how one console encodes this distinction. The Anchor Console presents Thailand's economic clusters through four scored axes, four geopolitical scenarios and four conditional review switches. We audit a fixed software revision and enumerate every binary switch configuration, without treating the encoded judgements as validated economic measurements. Four switches produce 16 distinct system states and 176 cluster-state rows: 32 rows differ from baseline and 144 do not. Across these rows, eight growth fields, sixteen sovereignty fields and sixteen categorical stances change; value and local-contribution scores remain fixed. The transitions are local to four different clusters, and repeated reads from the baseline do not accumulate upgrades. These are properties of the implementation, not estimated policy effects. Selected provenance checks also identify differences between encoded statistics and a later World Bank API vintage, a source-page mismatch, and console-derived geography. The full underlying foresight reports were unavailable for independent reading, so even the attribution of rule directions to those reports remains an encoded attribution. The contribution is an inspectable method for separating source claims, analyst conventions and executable consequences. The audit establishes neither forecast skill nor improved allocation, sovereignty, adoption or welfare.

Keywords: strategic foresight; Thailand; decision support; provenance; ordinal scorecards; scenario analysis; reproducibility.

1. Introduction

A policy scorecard can be easy to read and difficult to interrogate. A mark labelled High may summarise evidence, express a judgement or implement a display convention. Moving its plotted position after a switch is selected adds another claim: the reader now sees a conditional consequence. Without a record of how that movement was produced, an illustration can be mistaken for a forecast.

The Anchor Console is a useful case because it contains explicit distinctions worth testing. Its source register separates official statistics, secondary transcription, estimates, projections, judgements, derived material and external interpretation. Its threshold interface says that selecting a condition asks what would change if the condition held; it does not assert that the condition has occurred. Its scoring code identifies the size of an upgrade as a console convention. These disclosures make a bounded audit possible [1-3].

The console attributes its strategic argument to Poon Thiengburanathum's *The Anchor Strategy v1.1* and its economic scorecard to *Thailand's Economic Clusters V2*. It credits the console to Non Arkara/Axiom. We evaluate the encoded instrument, not a substitute for the unavailable reports [1,4].

Our question is narrow: **what consequences follow from the console's encoded review rules, and which boundaries must remain visible before those consequences can be interpreted as evidence?** We answer through source inspection, exhaustive state enumeration and selected primary-source checks. We do not ask whether the recommended strategy is optimal. Nor do we assume that a reproducible computation establishes the validity of its premises.

The contribution has three parts. First, we provide a complete, reproducible account of the four-switch mechanism. Second, we distinguish observations about software from economic interpretations, source attributions and future proposals. Third, we identify concrete requirements for a future governed scorecard: preserved source vintages, explicit scoring conventions and separately recorded evidence of institutional use.

2. Case, Materials and Scope

The deployed site is *The Anchor Console — Thailand in the Multipolar Transition*. The analytical object is repository revision 459d79c585d8c67830bd492ebe791429824f1024, not whichever revision is current when the paper is read. At inspection, the site's JavaScript and stylesheet were byte-identical to the local deployment assets. This verifies deployment identity at that capture, not the accuracy of every displayed statement [1].

The encoded scorecard has eleven economic clusters: food and agriculture; tourism; conventional automotive; EVs and batteries; electronics; data centres; petrochemicals; steel; finance; medical and wellness; and textiles. Each has four scored axes: economic value, growth, local contribution and sovereignty. A separate stance is a categorical recommendation, not a fifth scored axis. The source file expressly describes the scores as analytical judgements [2].

The strategic layer contains four named world trajectories and a six-row stress table covering five pillars and a sovereign foundation. The trajectories are presented as scenarios rather than calibrated forecasts. The threshold engine is a different object: it accepts a baseline cluster, its identifier and active review switches. It accepts no trajectory argument. The sixteen configurations studied here therefore must not be described as sixteen geopolitical futures or as a coupled geopolitical-economic simulation [3,5].

The geographic scope is Thailand, with national cluster arguments and regional presentation. Seven economic regions, eleven sectoral clusters and nine strategic geographic clusters are distinct classifications. Province membership in the nine strategic clusters is explicitly console-derived; the source notes that the underlying strategy names clusters without enumerating provinces. Enumeration of these memberships yields seventy unique provinces. That count describes encoded membership, not fieldwork coverage, policy adoption or empirical exposure. The map asset contains seventy-seven provinces [6].

There is an important source-access limit. The repository context states that the complete 139-page Anchor edition is absent and that its public Craft edition contains selected portions. The two Craft URLs returned document shells during retrieval, but complete readable report texts were not recovered. We therefore cannot independently verify report page references, score transcription, the elicitation process or the directions attributed to report section 6.1. References to these documents in this paper identify the console's stated lineage, not a full-text review [4].

3. Method

3.1 Fixed Inputs and Reproduction

The accompanying analysis.mjs reads eleven exact source files from the pinned revision. It computes SHA-256 hashes for the input bytes and executes the source modules in a temporary directory. It does not import the working tree's current scoring files. A source-snapshot mode supports the same file layout with a manifest declaring the commit and expected hashes. Both modes verify the input bytes before computation. The output records the source pin, input hashes, analysis hash and primary-check capture hash [7].

This distinction matters because the repository changed during preparation. A successful run against a newer working tree would not reproduce this case. The script also runs the existing test file against the pinned inputs. Passing tests are reported as implementation checks, not as independent validation of the source reports or their statistics.

3.2 Scales and Review Rules

The plotting scale encodes Very Low, Low, Low-Med, Med, Med-High and High as 0.5, 1, 1.5, 2, 2.5 and 3. Declining is a direction placed at the plotting floor; it is not substantively identical to Very Low. No cardinal welfare interpretation follows from equal distances on this display scale. We retain category labels in the results rather than aggregate them into an economic index [3].

The implementation defines four review switches. EV localisation raises the EV sovereignty score and removes its conditional stance. Data-centre governance raises sovereignty without removing conditionality. Steel restructuring changes its stance. Agricultural productivity raises food's growth category. The source code attributes these directions to the cluster report but explicitly supplies the magnitude: exactly one ordinal step where a scored axis changes. This audit verifies those encoded operations, not whether the inaccessible report states them faithfully or whether their conditions have been realised.

3.3 Exhaustive Enumeration

Four binary switches define 2^4 = 16 active sets. For each set, we evaluate all eleven clusters from their published-in-code baseline. A cluster-state row differs if any of its four scores or its stance differs. We count changed rows and changed properties separately because one row can change both a score and a stance. We retain all 176 rows in the output [7].

The checks compare computed changes with the accessor's reported moved fields, verify unchanged value and local contribution, and require clusters without an active targeting rule to retain their baseline profiles. They also compare reversed active-set insertion order and repeated reads from the same baseline. These are tests of this accessor and set representation. They are not tests of policy sequencing, and they do not imply that feeding an already adjusted profile into the lower-level function would be idempotent.

3.4 Selected Provenance Checks

We inspect the source register, geographic annotations, time-series module and declared pilot. We also retain selected World Bank API values retrieved on 5 November 2026 and compare them with the encoded 2024 values. The captured API metadata reports a last-update date of 8 October 2026, later than the audited software revision. The comparison can establish disagreement between two identified vintages, but cannot determine whether the original display was wrong when written [8].

These checks are purposive examples, not a representative statistical audit. Primary-source observations are preserved in primary-checks.json; arithmetic differences are regenerated by the script. No original report statistic is silently replaced by a later API value.

4. Results

4.1 State Space and Locality

All sixteen switch configurations produce distinct full-system profiles. There are 176 cluster-state rows, of which 32 differ from baseline and 144 do not. Each switch is active in eight configurations by construction. Because the four switches target four different clusters, each target contributes eight changed rows; the seven other clusters never change. There are fifteen distinct cluster-specific profiles across the eleven clusters: two for each target and one for each unaffected cluster [7].

Quantity	Result	Unit
Switch configurations	16	Exhaustively enumerated active sets
Cluster-state rows	176	Eleven clusters repeated across sixteen sets
Changed / unchanged rows	32 / 144	Comparisons with the same baseline
Growth changes	8	Category fields
Sovereignty changes	16	Category fields
Stance changes	16	Categorical recommendations
Value / local changes	0 / 0	Invariant fields

The forty changed properties are not forty changed rows. EV changes two properties each time its switch is active; the other targets each change one. Neither these counts nor their proportions are event frequencies, probabilities, treatment effects or empirical sample sizes.

4.2 The Four Transitions

Target	Before	After its switch is active
EV and batteries	Sovereignty Low; Focus (conditional)	Sovereignty Low-Med; Focus
Data centres and cloud	Sovereignty Very Low; Focus (conditional)	Sovereignty Low; stance remains conditional
Steel	Make it exist; growth Declining	Keep & upgrade; growth remains Declining
Food and agriculture	Growth Low-Med; sovereignty High	Growth Med; sovereignty remains High

The steel result is particularly instructive: a recommendation can change without any scored axis changing. Conversely, the data-centre score can improve without the recommendation becoming unconditional. A display that collapses these into one progress number would erase a distinction actually present in the rules. Food's switch changes the growth category but does not increase an already High sovereignty judgement.

No switch changes measured output, employment, ownership or fiscal expenditure. There is no estimated transmission coefficient from a rule activation to those quantities. The apparent precision of a plotted movement comes from the display encoding and the analyst's one-step convention.

4.3 Implementation Checks

The audit performs 176 locality checks, 176 insertion-order comparisons and 176 repeated-read comparisons. All pass. Baseline profiles remain unchanged when every switch is off, and the source cluster array is unchanged after enumeration. The pinned repository suite reports 25 tests, 25 passes and zero failures [7].

Repeated baseline reads do not ratchet an upgrade into a second upgrade. Reversing set insertion order produces the same output because this implementation evaluates its fixed rule list. These properties make the interface stable under repeated inspection, but do not establish how real policies interact over time. The enumeration contains no cross-cluster feedback, resource constraint or endogenous activation mechanism.

5. Provenance Findings

5.1 Source Presence Is Not Source Verification

The source register contains 55 records, and the literature module contains thirteen records. These are inventories. The tests check that references exist, tiers are declared and URLs have HTTPS syntax; they do not retrieve each source and establish that its content supports the attached claim. The time-series module is also outside the test helper's collected fact records [1,7].

Some news and aggregation links carry an official tier because they report an official body's figures. That preserves the claimed origin but obscures the transcription step. A research ledger should record the issuing agency and the intermediary separately. A functioning URL or a valid tier label cannot substitute for a specific table, observation date and extraction record.

5.2 Selected Statistical Disagreements

2024 indicator	Encoded value	Retrieved API value	Encoded minus retrieved
GDP per capita, current USD	7,900	7,386.636151	+513.363849 USD
Real GDP growth	2.5%	2.948600%	-0.448600 percentage points
FDI net inflows / GDP	1.1%	2.701653%	-1.601653 percentage points
Population aged 15-64 / total	70%	69.887763%	+0.112237 percentage points

The comparisons use the exact indicator identifiers named in the console. Differences are reproducible from the frozen selected responses, not a claim about values an API must return indefinitely. Original extraction vintages are unavailable. Revisions, rounding, transcription or definitional choices cannot be separated from this evidence alone; disagreement is not proof of fabrication or an original-vintage error [8].

The cited NXPO page provides a different kind of problem. Dated 23 June 2023, it reports 2021 R&D expenditure of 1.21% of GDP and the target of 2% by 2027. It does not establish the console's separately displayed 2023 value of 0.94%. This is a claim-to-page support gap, not evidence that 0.94% is false [9].

The console also encodes 78% as 2024 EV share, described as BEV plus PHEV. A BOI release reports 70,582 BEV passenger registrations in 2024 and a 15% share for new BEV passenger registrations in the first half of 2025. Different vehicle definitions, periods and sales/registration denominators prevent direct substitution. The combined-share claim requires its own primary series; the BOI figures are not a corrected replacement [10].

5.3 Geography and Added Interpretation

The national map shades regional GDP shares, not measured provincial cluster scores. The encoded Eastern region is labelled EEC in places although the strategic EEC membership lists three provinces. The broad regional denominator must not silently become a three-province corridor statistic. Likewise, province membership in a derived strategic cluster does not establish that every part of that province has the stated exposure [6].

The literature register assigns studies to regions through console-authored relevance tags. Counts derived from those tags describe a curated register, not the density of all research or a systematic evidence gap in a region. No search strategy or complete regional bibliography was established here. The map and register remain useful navigation instruments if those boundaries remain explicit.

6. Discussion

The strongest result is not that a particular sector should receive support. It is that the interface's consequences can be inspected without confusing a switch with an observation. The enumeration shows exactly which judgements move, which recommendation remains conditional, and which fields stay fixed. It exposes the small mechanism behind a visually rich instrument.

This also shows the limit of reproducibility. A correct implementation can consistently apply an uncalibrated convention. The one-step rule does not become an estimated effect because every state has been enumerated. The strategic layer's equal-weight stress summaries are not expected utilities under calibrated scenario probabilities; the reported alignment payoffs cannot be independently reconstructed without their underlying matrix. Normative dependence caps and spending envelopes should therefore remain proposals with stated assumptions [3,5].

The proposed operating system offers a sensible future research boundary. Its quarterly data spine, reviewed score records, threshold monitor and funding-decision surface are explicitly a proposed pilot, including a Northern Wellness and Creative Belt example. They are not completed deployments. A future study could record an analyst's evidence, a reviewer's score decision and the institution's allocation rationale. It would need to measure use separately from whether the resulting allocation improved an outcome [11].

The contribution here is a source-to-rule audit that another reader can challenge. It is not an automated substitute for judgement. Better provenance would allow the author to revise a premise, the analyst to identify the affected rule and the reader to see the consequence, without retroactively presenting a scenario as an observed success.

7. Limitations

This is one implementation audit, not an evaluation of Thailand's industrial policy. The underlying reports were unavailable in full; source fidelity and authors' elicitation methods remain unverified. The ordinal scores have no demonstrated inter-rater reliability or empirical calibration here. The configurations are constructed, equally enumerated possibilities, not equally probable futures. The checks omit real-world policy interactions and costs. Selected later-vintage API comparisons do not validate an entire series or adjudicate its historical correctness. There is no user experiment, adoption record, allocation-effect estimate or welfare outcome. Public reproduction additionally depends on release of the exact source supplement with appropriate rights; hashes alone are insufficient access.

8. Conclusion

An interactive foresight instrument can make its reasoning inspectable if observations, judgements and executable conventions remain distinct. At a fixed revision, four conditional switches generate sixteen states, changing only four of eleven clusters and leaving value and local contribution invariant. Those results establish the behaviour of the encoded instrument. They do not establish that its strategy works. The next step is governed evidence and recorded decisions, not a stronger empirical claim attached to the same animation.

Declarations And Reproducibility

Contributions. Research by Poon Thiengburanathum; analysis by Non Arkaraprasertkul, as directed by the project owner. Final author approval, detailed roles, affiliations, funding and competing interests remain to be confirmed. Neither institutional adoption nor institutional endorsement is implied.

AI assistance. An AI coding assistant prepared the audit script and this draft. The authors must review the analysis, claims and references before approving publication. Software checks and an independent provisional review are not scholarly peer review.

Materials. The methods supplement contains analysis.mjs, results.json, primary-checks.json, the manuscript and eleven pinned source files. Snapshot execution does not require the private Git repository. The supplement preserves the upstream Thailand geometry MIT notice separately; no blanket licence to third-party statistics, publications or source reports is asserted. The original reports are not redistributed.

Execution. With a Node version supporting TypeScript type stripping, run node --experimental-strip-types analysis.mjs --repo /path/to/Foresights --out reproduced.json, or use --snapshot /path/to/extracted-supplement. The pin and hashes are checked before execution. Compare the complete output with results.json. The script uses temporary copies, not a checkout or modification of the source repository. The primary checks are fixed audit captures, not live retrievals during reproduction.

References

1. The Anchor Console. Audited implementation, commit 459d79c585d8c67830bd492ebe791429824f1024: src/data/sources.ts, lines 17-62; src/lib/provenance.test.ts, lines 46-70 and 170-190. Live instrument: foresight.nonarkara.org. Source line references throughout refer to this commit.

2. Console cluster matrix. src/data/clusters.ts, lines 1-7, 23-287 and 300-312. Eleven baseline cluster records and baseline-resolving accessor.

3. Console scoring rules. src/lib/scoring.ts, lines 25-38, 101-140 and 157-200; src/components/Switchboard.tsx, lines 44-50. The latter is a contextual source locator, not one of the eleven execution inputs.

4. Thiengburanathum, P. The Anchor Strategy v1.1; Thailand's Economic Clusters V2 (July 2026). Console-attributed source documents: Anchor, Clusters. Full texts not recovered. Scope recorded in context.md, lines 9-20; no invented printed-page citations are supplied.

5. Console strategic layer. src/data/anchor.ts, lines 15-29, 46-102 and 200-242. Encoded scenarios, payoff constants and stress judgements, not independently verified forecasts.

6. Console geography. src/data/geography.ts, lines 4-15, 31-87 and 145-222; src/lib/thailandProvinces.ts, lines 40-66. Regional and derived strategic memberships.

7. Accompanying audit artifacts. analysis.mjs; results.json: counts, checks, singleSwitchEffects, states and inputHashes. These contain the complete enumerated results and pinned execution-input hashes.

8. World Bank. GDP per capita, GDP growth, FDI net inflows/GDP, working-age population share. Selected responses captured 5 November 2026; API last-update date 8 October 2026. Compared with src/data/thailandSeries.ts, lines 33-124; retained in primary-checks.json.

9. NXPO. Public policies and measures aim at boosting R&D expenditure to 2% GDP by 2027, 23 June 2023, opening expenditure paragraphs. Compared with src/data/corroboration.ts, lines 60-74.

10. Thailand Board of Investment. Press release 97/2568, 30 July 2025, physical PDF page 2 (unnumbered). BEV registrations and period-specific share. This is not the console's combined BEV/PHEV series.

11. Console pilot proposal. src/components/Pilot.tsx, lines 8-36, 65-71 and 132-135; src/components/System.tsx, lines 1-11. Proposed operating model, not evidence of execution.