METHODS PAPER

Net Zero Industrial Policy Lab · Johns Hopkins University

The Play Readiness Index

A subnational diagnostic for green industrial policy. Evidence, screening assumptions and the next questions for a location decision.

Coverage

32 states · three industrial opportunities

Version

v3.3 · September 2026

Research status

Revised working paper

00 Abstract

Industrial policy decisions require evidence on where productive capabilities coincide with conditions needed to support new investment. The Play Readiness Index combines export specialization, relatedness, employment concentration, recorded clean-capacity share, and industrial-park proximity to substations. An application to 32 Mexican states and three industrial plays compares observed profiles under explicit classification rules. Thirty-seven of 96 cells lack at least one scored input; replacing unobserved values with within-play medians changes 30 classifications. Weighting and aggregation choices also affect whether the canonical anchor retains its label. Sensitivity exercises identify dependence on these construction choices. Tests against subsequent investment, export, or employment outcomes remain necessary to assess predictive validity. The framework supports preliminary screening and the organization of further evidence. Facility decisions require additional information on electricity procurement, connection capacity, workforce availability, and commercial demand, gathered by developers, network institutions, and prospective buyers before a project proceeds.

01 The decision and the contribution

Export specialization, manufacturing employment, and electricity infrastructure describe different parts of an investment decision. A state can export a product intensively while offering limited evidence about the conditions facing an additional plant. State development agencies and investors need to bring these records together before commissioning site studies or allocating project-preparation resources. Comparing the underlying conditions provides a starting point for that investigation.

The unit of analysis is a state paired with an industrial play: a defined basket of products associated with a manufacturing opportunity. The Mexico application covers auto-supplier upgrading, electric-vehicle components, and grid hardware across all 32 federal entities, yielding 96 cells. These baskets define the application. Their selection does not establish that they exhaust Mexico's industrial opportunities, and individual products may have several end uses.

The contribution is an auditable combination of capability indicators and enabling conditions, followed by explicit classification rules. Balassa's specialization ratio measures a location's export composition relative to a reference economy [1]. Product-space research connects diversification to related productive capabilities [2], and regional evidence establishes that relatedness has a subnational application [3]. Economic complexity summarizes the structure of an economy's product network [4]; green complexity adapts that approach to environmentally relevant products [5]. The PRI adds a particular set of employment, energy-capacity, and infrastructure proxies to a state-by-play comparison. Its incremental value depends on the information those additions provide for a defined decision.

NZIPL's predictive-competitiveness framework estimates country-level export competitiveness across clean-energy technologies using supply-chain products, co-exports, and country characteristics [15]. The PRI addresses a subsequent geographic screening question with explicit rules and additional local proxies. It does not estimate that framework's predictive model at a smaller scale. The comparison identifies distinct analytical tasks and leaves their relative performance for a common outcome to future testing.

Framework	Object of comparison	Relationship to the PRI
Export specialization [1]	Product share relative to a reference economy	One input, adapted to a within-Mexico reference.
Regional relatedness [3]	Connections between existing and potential activities	Conceptual basis for the frozen HS4 approximation.
Economic and green complexity [4, 5]	Structure of broad or environmentally relevant export networks	Separate comparators with explicit product universes.
NZIPL predictive competitiveness [15]	Predicted country competitiveness in clean-energy technologies	A complementary national opportunity assessment.
PRI	Observed state-play profiles under stated screening assumptions	Auditable composition, rules, and sensitivity to local proxies.

The index and its categories have separate functions. A weighted score summarizes five recorded dimensions. An ordered set of rules then identifies profiles such as high specialization with low relatedness, or specialization alongside a low recorded clean-capacity share. These are conditional screening categories. Their usefulness requires evidence that they direct investigation toward consequential differences between locations. That question remains open until classifications are tested against outcomes beyond the inputs used to construct them.

Industrial-policy evaluation requires attention to objectives, implementation, and alternative explanations [6]. A category alone cannot establish which intervention will increase investment or welfare. Where a profile raises an electricity question, developers, Mexico's Federal Electricity Commission (CFE), and National Energy Control Center (CENACE) must establish the relevant supply and connection conditions; where it raises a capability question, firms and training providers must identify the actual workforce or supplier gap. The practical output is a documented question for those actors to resolve before an investment decision advances.

02 Energy conditions and policy scope

Installed capacity, electricity consumed by a plant, and the emissions attributed to its products are different measurements. Capacity is a stock measured in megawatts. Electricity generation and consumption are flows measured in megawatt-hours. Their relationship depends on operating hours, dispatch, trade between regions, and contractual arrangements. A state's recorded capacity share offers a geographic screening proxy with a narrower meaning than an electricity-consumption emissions factor.

The energy variable is the clean-capacity share recorded in the frozen PRODESEN-derived input. Its complement `F = 1 - C`, retained in the legacy `fossil_share` field, determines the author-defined energy screens. The input's clean category includes efficient cogeneration, biomass, and some storage alongside other technologies. These categories do not identify a zero-emissions or purely non-fossil share. The 0.60 anchor ceiling, 0.50 near-anchor ceiling, and 0.70 stranded-category floor apply to `F` as a recorded-share complement. No conversion to grams of CO2-equivalent per kilowatt-hour is assumed. Carbon variables preserved in the archived v3.2 comparison record play no role in the v3.3 score or classification.

Corporate Scope 2 accounting distinguishes electricity emissions estimated using a location-based approach from those attributed using qualifying contractual instruments. The GHG Protocol specifies data hierarchies and quality criteria for those methods [7]. A state boundary does not identify the grid serving every establishment, and a clean-power contract requires its own accounting and delivery assessment. Developers and electricity buyers should use the capacity screen to identify the records they need from suppliers and network operators before estimating a facility's emissions.

The EU Carbon Border Adjustment Mechanism applies to specified goods identified in the regulation's Annex I. Article 7 and Annexes II and IV govern which emissions enter the calculation and how they are measured; the treatment of indirect emissions depends on the covered product [8]. The implementation rules provide further calculation requirements, including conditions relevant to electricity emissions [9]. A customs-code and installation-level assessment is necessary for any product in a PRI basket. The capacity thresholds in the classifier confer no determination of CBAM liability, access, exemption, or compliance.

US clean-vehicle policy supplies historical context for changing supply-chain incentives. Section 30D requirements addressed vehicle and battery eligibility, including sourcing and assembly conditions [10]. The credit is unavailable for vehicles acquired after September 30, 2025; qualifying earlier acquisitions can have a later placed-in-service date [11]. This historical provision supplies no current basis for a universal state electricity-intensity screen. Exporters and their customs or tax specialists must check the relevant product, transaction date, and applicable provision when assessing market access.

A low recorded clean-capacity share raises an electricity-supply question while leaving the commercial feasibility of a particular plant unresolved. CFE and CENACE control the relevant network and system processes, and developers control project design and procurement choices within those processes. Project preparation should establish connection requirements, supply options, and the evidence requested by the buyer before attaching a financing recommendation to the profile.

03 Data, coverage, and measurement

Version 3.3 is a methodological revision dated September 6, 2026. It freezes the earlier three-play taxonomy and HS4 relatedness measure to keep changes in construction distinguishable from a change in the underlying product concept. It corrects state-name matching in employment, retains full computational precision in export rescaling, and replaces the former park score with the recorded substation-proximity measure. The research snapshot is separate from the continuously updated platform.

Input	Frozen observation period	Available cells	Measurement limit
OEC state-play exports and state totals	2025	69 of 96	Recorded product composition; attribution and coverage remain material.
INEGI state export totals	2024	All 32 states	Used to weight the reference composition; a different year and statistical universe.
HS4 relatedness approximation	2022	75 of 96	Frozen heading-based proximity, with no conversion to international product-space density.
DENUE-derived employment	May 2025, November 2024, and May 2024 releases	96 of 96	Broad, overlapping SCIAN mappings and mixed source dates.
PRODESEN installed capacity	December 2023, corrected June 2026	96 of 96	A state generation-capacity stock repeated across its three plays.
Recorded park proximity to substations	March 29, 2026 registry snapshot	60 of 96	A distance score with approximate coordinates and unverified connection capacity.

Fifty-nine cells have all five scored dimensions observed. The remaining 37 have at least one unobserved input. The registry contains 431 parks, but the state-play coverage of its proximity score is narrower than the full classification universe. These counts describe the frozen research inputs and remain separate from the live platform's current coverage.

Export specialization and the reference universe

Mexican state-product exports come from the stored OEC subnational extract. BACI is a separate country-to-country product-trade database [12]; it should not be described as the source of Mexican state observations. The INEGI state totals provide a second statistical universe. INEGI's Exportaciones Trimestrales por Entidad Federativa allocates exports using information about the location of producing establishments and reports quarterly and annual results [13]. Combining its totals with OEC product shares is an analytical rescaling whose assumptions remain visible.

Let `X_sp` denote OEC exports of play `p` from state `s`, `T_s` the OEC total for that state, and `I_s` its INEGI export total. The frozen OEC state data refer to 2025 and the INEGI totals to 2024. Method B retains the OEC composition `c_sp = X_sp / T_s` and calculates an INEGI-weighted reference share:

```
c_sp = X_sp / T_s
q_p  = sum_s(c_sp × I_s) / sum_s(I_s)
RCA_MX_sp = c_sp / q_p
allocated_exports_sp = c_sp × I_s
```

The numerator and reference share are computed at full precision. Rescaling changes the reference universe while preserving the recorded product composition within each state. Any product-level attribution error in that composition remains, and the mixed years prevent interpreting the allocated result as an observed annual export total. The allocated export value is a modelled quantity and is kept separate from the original OEC export value used by the canonical scale screen.

RCA-MX is a within-Mexico specialization ratio. A value of three means that the play occupies three times the national reference share of the state's recorded export basket. It does not establish productivity, domestic value added, the availability of local suppliers, or the size of a plant's expansion opportunity. Trade composition can also reflect ownership, firm concentration, or processing arrangements that the aggregate ratio cannot separate.

Relatedness and employment

The frozen relatedness measure uses the share of products within an HS4 heading that a state exports with RCA at least one, averaged across the products assigned to the play. Its scale runs from zero to one. HS4 membership supplies a reproducible approximation to productive proximity. It does not reproduce the international product-space proximity matrix, whose links are estimated from co-export patterns [2]. Adopting that alternative measure requires re-estimating thresholds and reporting the resulting classification changes.

The employment location quotient uses the DENUE national business directory through DataMexico. It compares the play-associated share of manufacturing employment in a state with the corresponding national share. The frozen SCIAN industry crosswalk uses subsector 336 for auto suppliers; 334, 335, and 336 for EV components; and 333 and 335 for grid hardware. These broad categories overlap and include activities outside the export baskets. The quotient measures employment concentration in the mapped activities. Occupation, skill, vacancies, wages, and commuting access require additional evidence.

State-name canonicalization occurs before every join. The legacy employment lookup failed to connect the Estado de México record to its canonical state key, assigning zero employment inputs to that state's cells. Version 3.3 restores those observations. Missing records, reported zeros, and unmatched keys have distinct meanings; a successful join is a prerequisite for interpreting the value that follows.

Generation capacity and park proximity

The PRODESEN input records state capacity at the end of 2023, drawing on the annex to the 2024–2038 electricity-system development program [17]. State capacity shares are retained in their original bounded scale. Their geographic aggregation cannot capture the generation dispatched to a particular load, network congestion, or the treatment of purchased electricity under a buyer's accounting requirements.

The park dimension is the state mean of the existing `grid_proximity` score for recorded industrial parks. This score derives from distance to a mapped substation. The legacy PRI actually consumed park `play_scores` derived from export specialization and employment, recycling capability inputs already represented elsewhere in the index. Its manuscript described a different composite involving proximity, clean energy, and play relevance. Version 3.3 aligns the implementation and definition around the proximity component alone, narrowing the construct to a physical-location proxy.

For each park, the inherited score is `round(max(0, 1 - (d / 200)^0.7), 3)`, where `d` is the recorded haversine distance in kilometres to the nearest mapped substation rated at least 230 kV. The stored distance has one decimal place. The state mean uses the resulting three-decimal park scores. All 431 inherited scores reproduce from their stored distances; the underlying geocoding and nearest-substation geometry are preserved from the registry snapshot.

A short mapped distance provides an initial geographic lead. It does not demonstrate spare substation capacity, an available connection, compatible voltage, or the cost and timing of required works. Some park locations are approximate or share coordinates, which further limits precision. Park operators and developers should obtain site coordinates and technical connection information from the relevant network institutions before treating proximity as an infrastructure advantage.

Incomplete observations

The canonical score assigns zero as the computational default for unobserved inputs so that all 96 cells remain traceable. A numerical default does not establish absence of exports, employees, parks, or related capabilities. Coverage flags accompany every classification, and sensitivity runs examine how alternative treatment of incomplete observations affects the result. Cells with complete records still carry measurement error; completeness indicates that an input was available, with no guarantee that it measures the intended construct closely.

04 Score and classification rules

Only RCA-MX and the employment location quotient receive within-play min-max normalization. RCA-MX is capped at 10 and the employment quotient at five before normalization. For a capped variable x, the normalized value is `(x - min(x)) / (max(x) - min(x))` across states in the same play; a constant column receives zero. Relatedness, clean-capacity share, and park proximity retain their bounded zero-to-one scales.

```
PRI_sp = 0.25 × RCA_norm_sp
       + 0.25 × relatedness_sp
       + 0.20 × employment_norm_sp
       + 0.15 × clean_capacity_share_s
       + 0.15 × park_proximity_s
```

The weights are author choices. They give export specialization and relatedness equal emphasis, employment a smaller share, and the two enabling-condition proxies equal shares. They estimate no structural production function. A high score can compensate for a low value on another dimension because the composite is additive; the classification rules subsequently impose selected non-compensatory conditions. The OECD and Joint Research Centre handbook provides the methodological rationale for making normalization, weighting, aggregation, and uncertainty explicit in composite indicators [14].

Within-play normalization makes scores sensitive to the comparison set. A state's normalized value can change when another state's observation changes. The scores support comparison within a fixed play and snapshot; comparing their levels across plays would require a common calibration. Percentile conditions use average ranks for ties, divided by the number of observations. Clean-capacity percentiles are calculated across the 32 states once, then attached to each play.

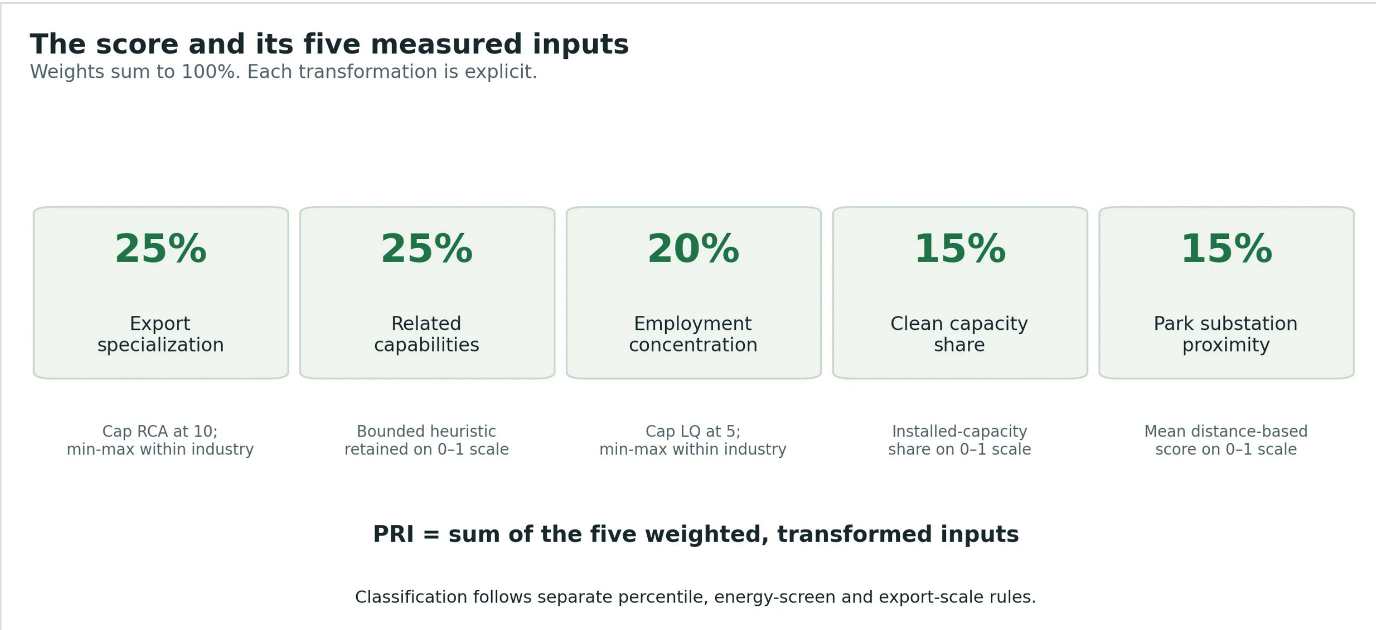

Figure 6. PRI construction. The informed score assigns 25% each to specialization and relatedness, 20% to employment, and 15% each to clean capacity and park substation proximity. Only capped RCA and employment use within-industry min-max scaling. The category rules apply after the score is calculated.

The classifier applies the following rules in order and assigns the first match. Here P is the percentile of the subscripted variable, R is RCA-MX, D is relatedness, C is recorded clean-capacity share, $F = 1 - C$ is its complement, and X is the original OEC play-export value in US dollars.

Order	Category	Conditions
1	Fragile exporter	`R ≥ 3` and `D ≤ 0.30`.
2	Anchor	`P_PRI ≥ 0.80`; `P_R`, `P_D`, `P_employment`, and `P_park ≥ 0.60`; `P_C ≥ 0.50`; `F < 0.60`; `X ≥ $300 million`.
3	Near-anchor	`P_PRI ≥ 0.70`; at least four of `P_R ≥ 0.50`, `P_D ≥ 0.50`, `P_employment ≥ 0.50`, `P_C ≥ 0.40`, and `P_park ≥ 0.50`; `F < 0.50`.
4	Stranded capacity	`F ≥ 0.70`, `P_R ≥ 0.60`, and `R ≥ 1`.
5	Peripheral	`P_R`, `P_employment`, and `P_park ≤ 0.30`.
6	Infrastructure orphan	`C ≥ 0.40`, `P_C ≥ 0.70`, `P_park ≤ 0.40`, `P_R ≤ 0.40`, and `P_employment > 0.30`.
7	Latent reserve	`0.40 ≤ P_PRI ≤ 0.70` and `P_C ≥ 0.60`.
8	Transitional	Every preceding rule is false.

The near-anchor checks are deliberately less demanding than the anchor checks. Its energy ceiling is nevertheless stricter, an inherited asymmetry retained as an explicit assumption and included in threshold sensitivity. A cell can also have a low PRI yet meet an earlier absolute rule. Publishing precedence alongside the thresholds makes these outcomes reproducible.

The $300 million floor applies to anchor classification using original OEC exports. It is an author-defined scale screen, with no claim to identify a universal minimum viable cluster. An alternative based on rescaled exports answers a different question and is tested separately. Holding the floor constant in an energy-guard ablation is necessary to attribute classification changes to the energy condition alone.

The category names preserve continuity with the earlier release. “Fragile exporter” identifies high specialization with low relatedness; its exposure to actual shocks remains an empirical question. “Stranded capacity” identifies an export and capacity-mix profile; no asset impairment, lost contract, or financing denial is observed by that rule. “Infrastructure orphan” describes a combination of recorded proxies. “Transitional” means that none of the preceding conditions holds, leaving the cause of that residual status open.

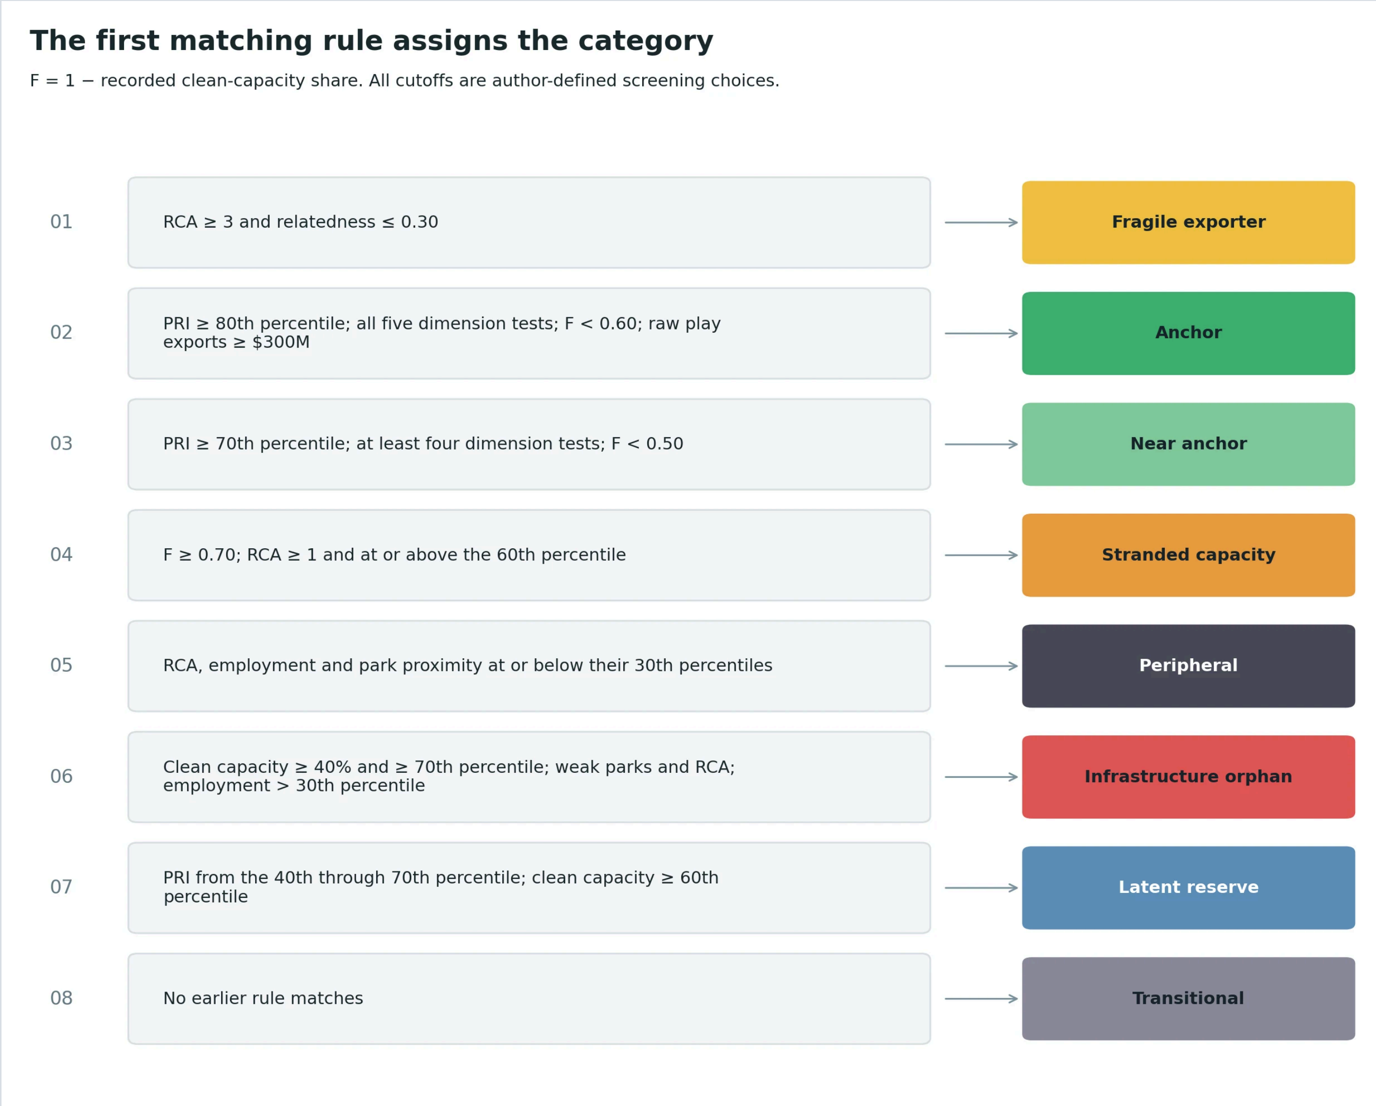

Figure 7. Classification rules. The classifier evaluates the rules in this order. The full dimension cutoffs and operators appear in the method table. Energy guards, the export-scale floor and the stranded-capacity rule can be varied independently. The terminal category records an unmatched pattern.

05 Mexico application and worked examples

The canonical run assigns one cell to anchor, five to near-anchor, and 18 to stranded capacity. Those totals describe the stated conjunctions of inputs and thresholds. They provide no direct count of investable locations or impaired assets. State agencies and developers should inspect the failed conditions and underlying records before using a category to define a project-preparation list.

Category	Auto suppliers	EV components	Grid hardware	Total
Anchor	1	0	0	1
Near-anchor	3	1	1	5
Fragile exporter	0	2	2	4
Stranded capacity	4	7	7	18
Peripheral	5	4	7	16
Infrastructure orphan	2	2	0	4
Latent reserve	2	4	4	10
Transitional	15	12	11	38
Total	32	32	32	96

Seven classifications change relative to the archived v3.2 output after the joint corrections and revised park definition. The anchor, fragile-exporter, stranded-capacity, infrastructure-orphan, and peripheral counts remain unchanged, while cells move among near-anchor, latent reserve, and transitional. This comparison describes the combined revision; assigning a change to one correction requires the corresponding isolated scenario.

Figure 1. State-industry classifications. The revised rules classify all 96 state-industry cells. 1 cell meets the anchor criteria and 18 meet the stranded-capacity pattern. Values compare locations within each industry. A dagger marks cells with unavailable inputs and computational defaults. A category identifies evidence to investigate; it does not verify a site or a market-access condition.

Aguascalientes: reconstructing an anchor assignment

Aguascalientes/auto-supplier has recorded play exports of $1.351 billion within an OEC state total of $2.770 billion. Its composition is 0.48774 and its Method B RCA-MX is 3.46539. The raw relatedness value of 0.333 exceeds the fragile-exporter ceiling of 0.30, so the first rule does not match. Its five dimension percentiles pass the anchor tests, the complement of recorded clean share is zero, and original exports exceed $300 million.

Dimension	Raw value	Scored value	Weight	Contribution
RCA-MX	3.46539	0.90005	0.25	0.22501
Relatedness	0.33300	0.33300	0.25	0.08325
Employment quotient	2.70000	0.85669	0.20	0.17134
Clean-capacity share	1.00000	1.00000	0.15	0.15000
Park proximity	0.91500	0.91500	0.15	0.13725
PRI			1.00	0.76685

The anchor assignment follows from this arithmetic and the published tests. A clean-capacity share of one still leaves the plant's actual electricity supply unresolved. Developers and electricity buyers should verify supply contracts and the proposed load; CFE and CENACE must establish the applicable network conditions before the classification informs a site commitment.

Zacatecas: identifying more than one failed condition

Zacatecas/auto-supplier combines RCA-MX of 3.85021 with relatedness of 0.333 and PRI of 0.73327. It fails the anchor park-percentile condition: its proximity score of 0.877 places it at percentile 0.46875, below the required 0.60. Original OEC play exports of $68.94 million also fall below the scale floor. Four of the five lower near-anchor checks pass and the energy ceiling passes, so the third rule assigns near-anchor.

Method B allocates $2.142 billion of exports to this cell, illustrating the difference between the original and rescaled quantities. Substituting that quantity would clear the scale screen while leaving the failed park condition. Analysts evaluating the cell should retain both quantities and identify every failed test; developers can then seek better park and connection evidence before interpreting the difference as an actionable infrastructure gap.

Tlaxcala/EV-components illustrates a different rule. RCA-MX of 7.43458 and relatedness of 0.030 trigger fragile exporter before the PRI is considered. Its score of 0.57944 cannot override that assignment. Firms and state development agencies should examine product, customer, and supplier concentration to determine whether the recorded specialization is exposed to a specific shock before designing a diversification response.

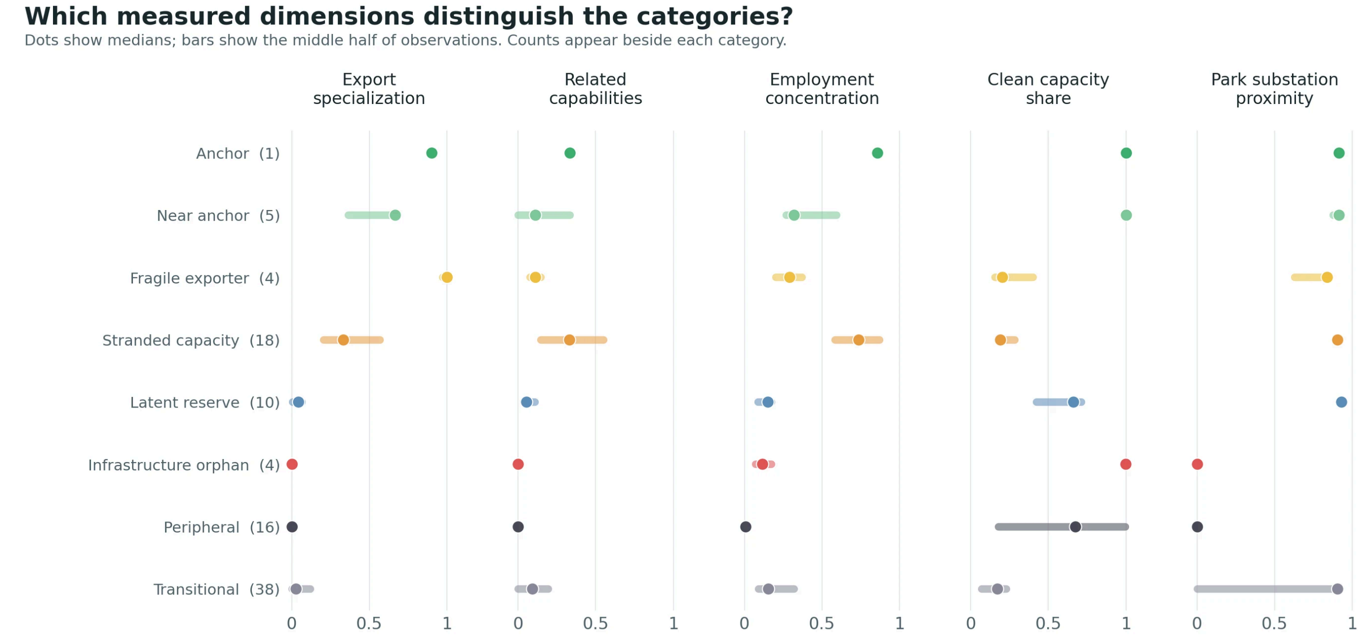

Figure 2. Dimension profiles. Category medians and interquartile ranges use the exact five model inputs. RCA and employment are capped and min-max scaled within industry; the other inputs retain their bounded scales. The park measure describes substation proximity. Computational defaults for unavailable inputs remain included. Categories with no cells have no plotted profile.

06 Sensitivity and benchmark comparisons

Weight perturbation, rule sensitivity, and comparison with another index address different questions. Weight perturbation asks whether a chosen profile survives alternative emphasis on the same inputs. Rule sensitivity changes the boundaries of the categories. Benchmark comparison shows how rankings differ under another construct. None supplies an observed investment outcome against which the PRI's usefulness can yet be assessed.

The 18-specification grid crosses equal weights of 0.20, canonical weights, and PCA-derived weights with arithmetic and geometric score aggregation and mean, median, and maximum park aggregation. PCA weights are squared loadings of the first principal component of standardized normalized dimensions, estimated across all 96 cells and refitted for each park aggregation. The geometric specification is `exp(sum_j(w_j × log(x_j + 0.01))) - 0.01`; the offset permits zero-valued inputs. Random weight perturbations add independent uniform draws from −0.10 to 0.10, floor the resulting weights at 0.05, and renormalize them to sum to one. This is Monte Carlo sensitivity. Since observations are not resampled, assignment frequencies should not be read as sampling confidence intervals.

Eighty-nine cells receive the same category across all 18 weighting and aggregation specifications. Aguascalientes/auto-supplier is an anchor in 12 specifications and a near-anchor in six, which limits the stability of the canonical one-anchor result. Across 1,000 weight perturbations of up to ±0.10, followed by renormalization with seed 42, 85 cells receive their modal category in at least 90% of draws. Aguascalientes/auto remains an anchor in all those draws. The different exercises explore different neighborhoods of the model.

The threshold grid crosses five sets of three values, yielding 243 scenarios: relatedness ceilings of 0.25, 0.30, and 0.35; stranded-category floors on F of 0.65, 0.70, and 0.75; anchor ceilings on F of 0.50, 0.60, and 0.70; near-anchor ceilings on F of 0.40, 0.50, and 0.60; and scale floors of $100 million, $300 million, and $500 million. RCA's fragile threshold remains three. Relative to the canonical run, scenarios change between zero and nine cells, produce zero or one anchor, and identify between 11 and 18 stranded-capacity cells. Those ranges describe a chosen set of assumptions, with no probability assigned to each specification.

Raising the relatedness ceiling to 0.35 alone reclassifies Aguascalientes/auto and Zacatecas/auto as fragile exporters. Raising the stranded-category floor on F to 0.75 alone reclassifies seven cells and leaves 11 in that category. The associated category totals depend on where these two boundaries are placed.

Incomplete observations produce larger changes than several weighting choices. Replacing unobserved inputs with observed within-play medians changes 30 of 96 classifications. Restricting the estimation to the 59 complete cells and refitting normalization and percentiles changes 16 of those 59 categories; that exercise changes both coverage and the comparison population. Renormalizing score weights over available components changes four cells while retaining the baseline treatment of raw inputs in classification rules.

Component correlations also depend on coverage. The normalized RCA component and park proximity have a Spearman correlation of 0.542 across all 96 cells with default values, compared with −0.453 across 59 jointly observed pairs. The sample and treatment of absent inputs change together. Removing mechanical reuse of capability from the park definition leaves statistical relationships dependent on the observations and specification used to estimate them.

Omitting one component from the composite, then renormalizing the remaining weights while retaining all raw-variable rules, changes between one and six categories. Replacing mean park proximity with the median changes none; using the maximum changes three. Removing the 93 municipality-centroid park locations also changes none. These checks leave the other approximate coordinates and the absence of verified connection capacity unresolved.

A hypothetical measurement stress shifts recorded clean-capacity share in every state by five percentage points, bounds the result to zero through one, and recomputes the score and rules. Subtracting five points changes no category; adding five changes seven. These shifts are chosen stresses and have no interpretation as estimated measurement confidence bounds. A separate diagnostic reconstructs clean-capacity shares from the frozen technology table using either its own total or the retained aggregate total; both produce zero classification changes. Verification against updated primary plant records would require a separate data revision.

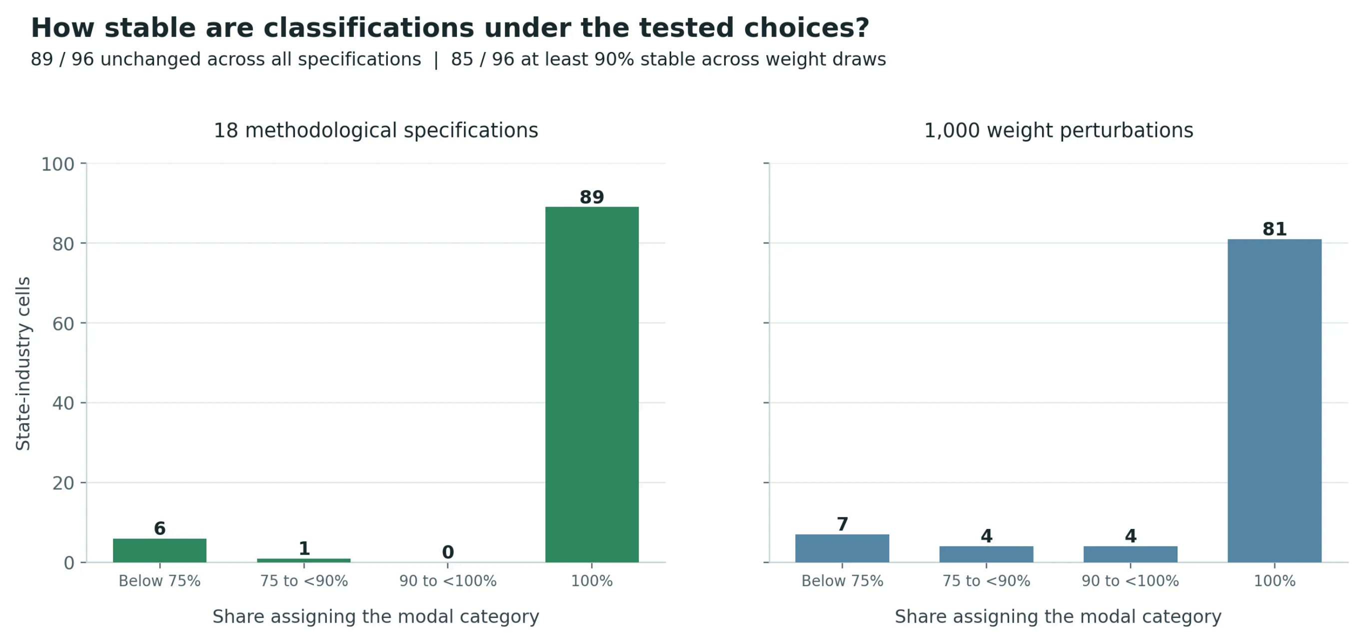

Figure 3. Classification sensitivity. All 96 observations appear in each histogram, including values equal to 1. 89 cells retain one category across 18 specifications; 85 have a modal category in at least 90% of 1,000 weight draws. Threshold and measurement sensitivity are reported separately. These exercises assess model behavior.

An isolated energy-guard counterfactual changes only the anchor and near-anchor checks on F. All other rules, including the scale floor and stranded-capacity category, remain active. A separate taxonomy experiment can remove the stranded-capacity branch, but its effects have a broader interpretation. The former comparison estimates dependence on two guards; the latter changes what kinds of profiles the taxonomy can recognize.

Removing only the two energy guards changes 19 classifications. Fifteen stranded-capacity cells move to anchor or near-anchor and four transitional cells move to near-anchor; three stranded-capacity cells retain that label. Removing the stranded-capacity branch alone changes 18 labels by construction. Removing both the guards and that branch changes 22 cells. Each count refers to its specific intervention in the classifier.

Removing the scale floor alone changes no canonical category, as does substituting Method B allocated exports for the original OEC scale quantity. Other failed conditions prevent a promotion in this snapshot. These null changes narrow the interpretation of the scale rule's current effect while leaving its influence under other inputs or thresholds open to further testing.

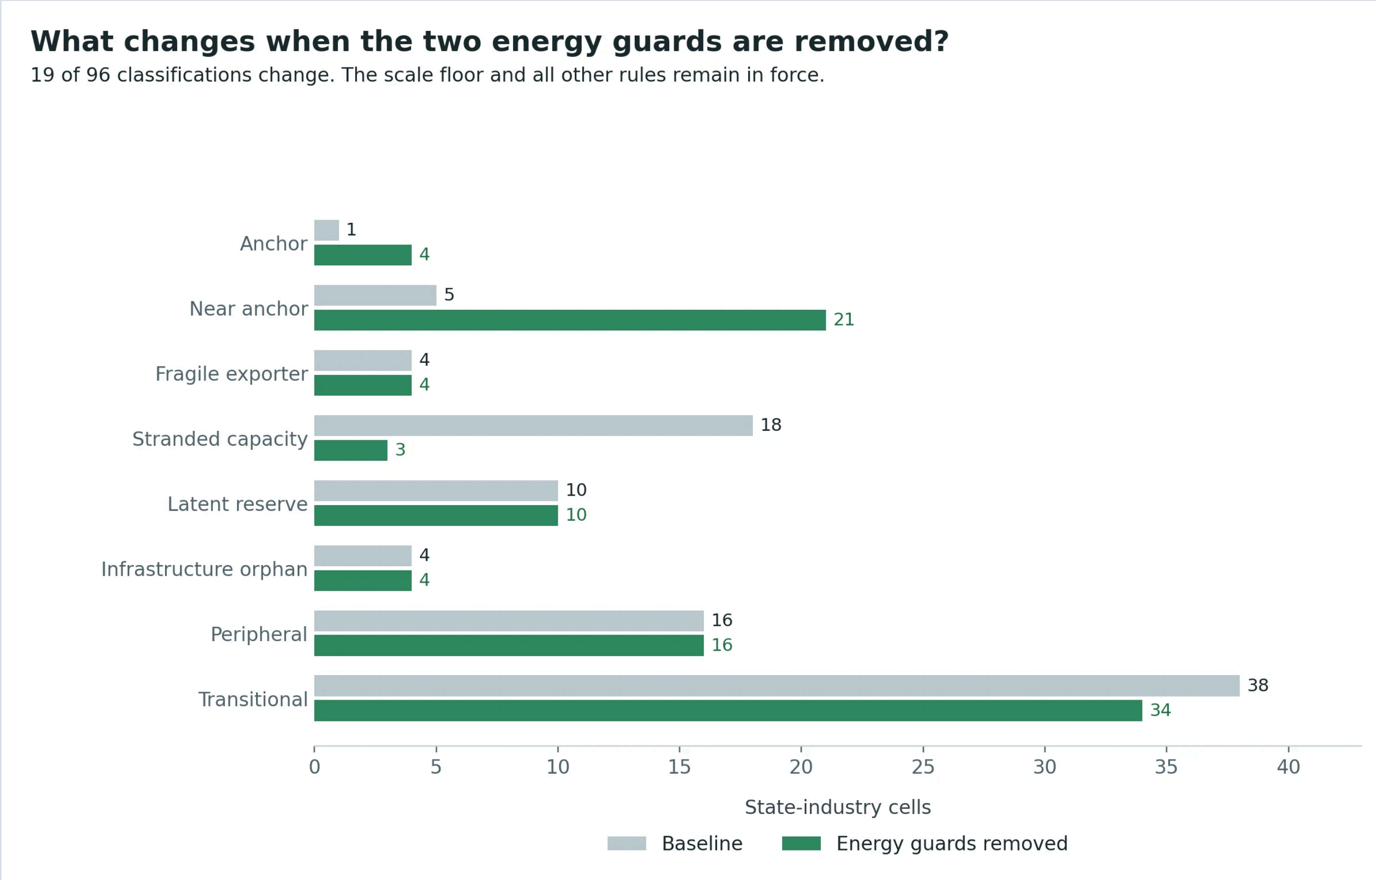

Figure 5. Isolated energy-screen comparison. Removing only the guards on the complement of recorded clean capacity changes 19 of 96 classifications. 15 cells move from stranded capacity to anchor or near-anchor. The export-scale floor, stranded-capacity rule, continuous score and other conditions remain fixed. This is a model sensitivity experiment.

The broad economic-complexity comparator uses the available DataMexico cache of 32 states and 1,961 products pooled across 2024 and 2025. Its wider product universe corrects the earlier comparator's restriction to selected play products. The pooled vintage differs from the canonical PRI export snapshot, so the comparison cannot isolate differences in method from changes in timing. The matrix treats 53,768 absent state-product cells as zero; reporting gaps and structural zeros cannot be distinguished, and the cache does not record the HS revision. An additional historical comparator uses a 2024 matrix of 25 states and 36 selected products.

Within each matrix, state-product RCA compares a state's product share with the share across reported states. A binary matrix `M` records RCA at least one, with row sums `d` measuring diversity and column sums `u` measuring ubiquity. ECI is derived from the second right eigenvector of `diag(d)^-1 M diag(u)^-1 M′`, calculated through its symmetric equivalent [4, 16]. The vector is oriented to correlate positively with diversity and standardized using the sample standard deviation across that matrix's states. Independent iteration checks the eigenvector calculation. The restricted comparator's diversity correlation is only 0.037, making the substantive interpretation of its sign weak.

All pairwise correlations within a play use the same complete-observation sample. A cell is included only when all five PRI inputs are observed and every comparator is defined. This yields 19 auto-supplier states, 20 EV-component states, and 20 grid-hardware states. Unreported-state exports of $3.575 million are excluded from every denominator in the broad matrix; its $591.462 billion of reported-state exports reconcile to the stored national cache, without establishing that the extract exhausts actual exports.

Play	Common states	PRI versus RCA-MX	PRI versus broad ECI	PRI versus restricted historical ECI
Auto suppliers	19	0.928	−0.165	0.751
EV components	20	0.798	0.468	−0.123
Grid hardware	20	0.815	0.365	−0.459

The coefficients are descriptive Spearman rank correlations. Across all three plays, PRI aligns more closely with RCA-MX than with broad ECI, while the direction of the broad ECI association varies by play. The coefficients reflect the joint choice of product universe, source vintage, and common sample. No confidence interval, significance test, or predictive performance claim is attached to these deterministic snapshot comparisons.

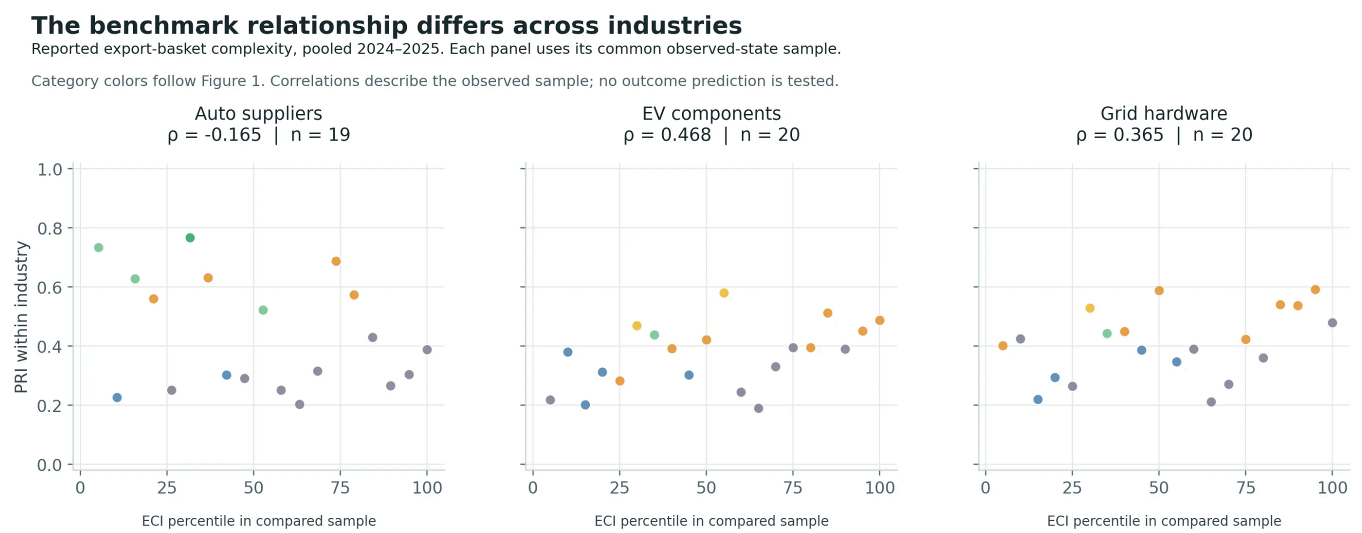

Figure 4. PRI and reported export-basket complexity. Auto suppliers: ρ = -0.165, n = 19; EV components: ρ = 0.468, n = 20; Grid hardware: ρ = 0.365, n = 20. The ECI uses 1,961 reported HS6 products across 32 states. The horizontal axis shows ECI percentiles within each panel's common observed-state sample. Pooled trade vintages, missing records and sparse state baskets limit interpretation. No regression line or causal effect is estimated.

A correlation with ECI measures agreement between two rankings. A negative value indicates that one tends to rank highly the states the other ranks lower; a positive value indicates greater alignment. Neither direction establishes which ranking better predicts an outcome. The published Green Complexity Index aggregates product-complexity information across competitive green exports [5]. Its reproduction requires an audited green-goods mapping and suitable product-complexity estimates, so a GCI result is excluded from this release.

07 External validation protocol

Predictive validity remains untested. A prospective evaluation should freeze the taxonomy, input vintage, transformations, and thresholds before collecting outcome data. The primary outcomes should be specified separately: entry into the play's export basket, sustained export growth, manufacturing employment growth, and independently verified establishment or expansion of a facility. Each measures a different aspect of readiness, and the observation window should reflect the time required for that outcome to occur.

A time-separated design would construct the PRI from information available at a baseline date and evaluate outcomes over a subsequent period. Baseline exports, employment, and population should enter simple comparison models before the composite is added. RCA-MX, relatedness, and broad ECI should each receive a comparable test on the same state-play sample. Out-of-sample error, calibration where probabilities are estimated, and the performance of a short list of candidate locations would make the practical cost of classification visible.

Repeated observations for the same state share infrastructure and institutions. Evaluation should account for that dependence, document the small number of independent geographic units, and avoid treating 96 cells as 96 unrelated observations. Broader time coverage or additional countries could improve the design, provided the measurement and classification rules are harmonized before the comparison. Threshold selection and performance assessment should use separate data partitions to limit tuning to the observed outcomes.

The fragile-exporter hypothesis requires its own test. Product concentration, buyer or firm concentration where observable, and historical export volatility should be measured alongside relatedness. A documented product or customer shock would permit comparison of subsequent adjustment among similarly specialized cells with different relatedness profiles. Such a design could establish whether the conjunction provides information beyond the simpler concentration measures; the present cross-section cannot estimate that relationship.

Testing an energy constraint requires evidence about electricity actually available to the affected facility. Connection delays, interruptions, procurement contracts, and verified emissions requirements could distinguish several mechanisms that a state capacity share combines. CFE, CENACE, plant operators, and buyers hold different parts of that record. A research partnership should specify those fields, their dates, and access conditions before interpreting a classification change as relief of a binding constraint.

08 Use, limits, and reproducibility

The index organizes preliminary investigation at a geographic scale where several relevant records can be joined. A favorable profile can justify further project preparation while leaving delivery conditions unresolved. State agencies and developers should carry the raw inputs, coverage flags, and sensitivity results into the site assessment, then record which additional evidence changed the decision.

An export-intensive cell with a low recorded clean-capacity share raises a supply question. The relevant action depends on the plant's actual connection and procurement options. CFE and CENACE must establish network conditions, developers must assess project design and contracts, and lenders must evaluate the resulting cash flows and risks. The PRI can direct attention to that sequence; an instrument recommendation requires the additional evidence it produces.

A high-specialization, low-relatedness profile raises a capability question. Supplier and workforce evidence can establish whether nearby production supports diversification or whether the export position depends on a narrow set of firms and customers. Firms, training institutions, and state development agencies can use that information to design a targeted preparation program. Development banks should assess the specific proposed activity and its financing need before assigning a credit, guarantee, or grant structure.

The principal limits are geographic aggregation, mixed vintages, broad employment concordances, approximate park coordinates, incomplete records, and author-defined thresholds. The generation mix and park proximity are repeated across plays within a state, while some mapped activities and products overlap. These shared inputs constrain interpretation of apparent cross-play agreement. The framework's portability depends on comparable subnational records and a fresh assessment of its thresholds in the new setting.

The replication package freezes inputs separately from platform feeds and records their hashes, source descriptions, transformations, and data vintages. It publishes raw cells with nulls and observation flags, the 96-cell scored table, classification parameters, sensitivity specifications, and benchmark construction. Regeneration uses those local inputs, with random seeds recorded for simulated perturbations. A reviewer can trace a displayed category to its raw observations, computational defaults, normalized values, percentile tests, and first matching rule.

The replication package contains the frozen data, source manifests, scripts, and pinned dependency specification. From its extracted root, create an environment and regenerate the analytical outputs:

```
python3 -m venv .venv
.venv/bin/python -m pip install -r research/pri-v3.3/requirements.txt
.venv/bin/python research/pri-v3.3/pipeline.py
.venv/bin/python research/pri-v3.3/benchmark.py
.venv/bin/python research/pri-v3.3/figures.py
```

Input hashes are checked before the core calculation. The analytical scripts use local snapshots and require no network requests. Figure regeneration follows the numerical outputs, so changes to the model can be checked against the tables and diagrams that communicate its results.

09 References

1. Balassa, B. (1965). Trade Liberalisation and "Revealed" Comparative Advantage. *The Manchester School*, 33(2), 99–123. doi:10.1111/j.1467-9957.1965.tb00050.x.

2. Hidalgo, C. A., Klinger, B., Barabási, A.-L., and Hausmann, R. (2007). The Product Space Conditions the Development of Nations. *Science*, 317(5837), 482–487. doi:10.1126/science.1144581.

3. Neffke, F., Henning, M., and Boschma, R. (2011). How Do Regions Diversify over Time? Industry Relatedness and the Development of New Growth Paths in Regions. *Economic Geography*, 87(3), 237–265. doi:10.1111/j.1944-8287.2011.01121.x.

4. Hidalgo, C. A., and Hausmann, R. (2009). The Building Blocks of Economic Complexity. *Proceedings of the National Academy of Sciences*, 106(26), 10570–10575. doi:10.1073/pnas.0900943106.

5. Mealy, P., and Teytelboym, A. (2022). Economic Complexity and the Green Economy. *Research Policy*, 51(8), 103948. doi:10.1016/j.respol.2020.103948.

6. Juhász, R., Lane, N., and Rodrik, D. (2024). The New Economics of Industrial Policy. *Annual Review of Economics*, 16, 213–242. doi:10.1146/annurev-economics-081023-024638.

7. World Resources Institute and World Business Council for Sustainable Development. (2015). *GHG Protocol Scope 2 Guidance: An Amendment to the GHG Protocol Corporate Standard*. Official guidance.

8. European Parliament and Council. (2023, consolidated October 20, 2025). Regulation (EU) 2023/956 establishing a carbon border adjustment mechanism. Article 7 and Annexes I, II, and IV. EUR-Lex consolidated text.

9. European Commission. (2025). Commission Implementing Regulation (EU) 2025/2547 laying down rules for calculating emissions embedded in goods. Official regulation.

10. US Department of the Treasury and Internal Revenue Service. (2024). Clean Vehicle Credits under Sections 25E and 30D; Transfer of Credits; Critical Minerals and Battery Components; Foreign Entities of Concern. Treasury Decision 9995, *Internal Revenue Bulletin* 2024-22. Official final rules.

11. Internal Revenue Service. (2026). *Clean Vehicle Tax Credits*. Acquisition cutoff and placed-in-service guidance. Accessed September 6, 2026. Official guidance.

12. Gaulier, G., and Zignago, S. (2010). *BACI: International Trade Database at the Product-Level. The 1994–2007 Version*. CEPII Working Paper 2010-23. Original methodology; current dataset documentation.

13. Instituto Nacional de Estadística y Geografía. (2020). *Síntesis Metodológica de las Exportaciones Trimestrales por Entidad Federativa*. Official program and methodology.

14. OECD, European Union, and European Commission Joint Research Centre. (2008). *Handbook on Constructing Composite Indicators: Methodology and User Guide*. OECD Publishing. doi:10.1787/9789264043466-en.

15. Ratan, I., Goldman, J., Han, R., Bagozzi, B. E., Sahay, T., and Allan, B. (2025). *Mapping the Industrial Base for the New Energy Economy: Predicting Competitiveness in Clean Energy Value Chains*. NZIPL Working Paper No. 4, November version. Institutional working paper.

16. Mealy, P., Farmer, J. D., and Teytelboym, A. (2019). Interpreting Economic Complexity. *Science Advances*, 5(1), eaau1705. doi:10.1126/sciadv.aau1705.

17. Secretaría de Energía. (2024). *Programa de Desarrollo del Sistema Eléctrico Nacional 2024–2038*. Annex I. Official capacity tables.

HOW TO CITE

García-Vazquez, G. The Play Readiness Index: a subnational diagnostic for green industrial policy. NZIPL working paper v3.3, September 2026. Net Zero Industrial Policy Lab, Johns Hopkins University. https://lineupmx.com/methods-paper

Net Zero Industrial Policy Lab · Johns Hopkins University | Mexico · v3.3 · netzeropolicylab.com