Beyond megawatts: Structural configurations and project ecologies in global utility-scale solar

Hairong Wang*[1], Chengyi Hong[1], and Jun Sun†[1]

[1]Civil and Environmental Engineering, Carnegie Mellon University, Pittsburgh, USA

1 Supplementary Information

SI.1 Data and retained scope details

SI.1.1 Source table and retained statuses

The analytical base uses the February 2026 release of the Global Solar Power Tracker and is restricted to the utility-scale main sheet (*Utility-Scale (1 MW+)*), interpreted at the phase level. Distributed-generation material was excluded because the main paper compares national utility-scale project structure rather than mixed utility-scale and distributed-generation records. Four statuses were retained in the main analysis: *operating*, *construction*, *pre-construction*, and *announced*. Together, these statuses define the active stock-plus-pipeline frame used throughout the main paper.

SI.1.2 Retained sample construction

Table S1 summarizes the retained sample and baseline analytical scope. After status and capacity filtering, the source table contained 213 countries and 100,028 phase-level rows. Applying the predefined country threshold of record count ≥ 30 produced a retained analytical base of 94 countries, 98,942 rows, and 3,573,038.8 MW. The threshold was used as a profile-stability rule rather than a size-weight rule: the purpose was to retain countries with sufficient phase-level observations to support a stable size–status profile, rather than to privilege countries with the largest absolute installed or pipeline capacity.

The predefined size bins were 1–5, 5–50, 50–500, and 500+ MW, with boundary rules of $1 \leq \text{MW} < 5$, $5 \leq \text{MW} < 50$, $50 \leq \text{MW} < 500$, and $\text{MW} \geq 500$. Combined with the four retained statuses, these bins yielded 16 predefined size–status modes. Before the country threshold was applied, the retained status composition was dominated by operating rows (81,377 rows; 1,267,798.9 MW; 198 countries), followed by pre-construction (12,715 rows; 1,143,626.7 MW; 135 countries), announced (3,092 rows; 785,275.0 MW; 153 countries), and construction (2,844 rows; 469,310.7 MW; 114 countries). China was the largest retained country by record count before country-threshold application, with 17,320 rows and 1,429,088.6 MW.

SI.1.3 Country filter and bin assignment audit

Table S2 reports the country-threshold edge cases and the bin-assignment audit. The threshold behaved transparently at the boundary: Mauritania and Puerto Rico were retained at exactly 30 rows, while Ecuador and Madagascar fell just below the threshold and were excluded. This confirms that the country filter was applied deterministically and that threshold-edge cases are directly inspectable.

The bin-assignment audit confirms that retained rows were assigned cleanly across the four predefined size bins. The 1–5 MW bin contained 55,743 rows and 107,413.2 MW, the 5–50 MW bin contained 28,040 rows and 478,419.5 MW, the 50–500 MW bin contained 14,131 rows and 1,758,342.1 MW, and the 500+

*Corresponding author: hairongw@andrew.cmu.edu
†Corresponding author: juns2@andrew.cmu.edu

MW bin contained 1,028 rows and 1,228,864.0 MW. All 16 size–status modes remained assigned in the retained analytical base. The retained scope used in the main paper is therefore substantively bounded and audit-traceable.

Table S1: Retained sample and analytical scope summary.

Item	**Value**
Panel A. Source table and baseline scope	
Dataset release	Global-Solar-Power-Tracker-February-2026.xlsx
Source sheet	Utility-Scale (1 MW+)
Source scope	Utility-scale main sheet only
Observational unit	Phase-level
Distributed sheet excluded	True
Retained statuses	Operating; construction; pre-construction; announced
Excluded statuses	Cancelled; shelved – inferred 2 y; cancelled – inferred 4 y; shelved; retired; mothballed
Country filter rule	Record count ≥ 30
Size bins	1–5; 5–50; 50–500; 500+
Bin boundary rule	$1 \leq \text{MW} < 5$; $5 \leq \text{MW} < 50$; $50 \leq \text{MW} < 500$; $\text{MW} \geq 500$
Panel B. Retained sample totals	
Countries before country filter	213
Countries retained	94
Countries removed by country filter	119
Rows after status and capacity filters	100,028
Rows retained	98,942
Rows removed by country filter	1,086
Retained capacity (MW)	3,573,038.8
Capacity removed by country filter (MW)	92,972.5
Minimum retained country record count	30
All 16 modes assigned	True
Panel C. Retained status composition before country threshold	
Operating rows / MW / countries	81,377 rows; 1,267,798.9 MW; 198 countries
Construction rows / MW / countries	2,844 rows; 469,310.7 MW; 114 countries
Pre-construction rows / MW / countries	12,715 rows; 1,143,626.7 MW; 135 countries
Announced rows / MW / countries	3,092 rows; 785,275.0 MW; 153 countries
Largest retained country by record count	China (17,320 rows; 1,429,088.6 MW)

The retained analytical scope was fixed before country projection, community detection, and sensitivity analysis.

SI.2 Robustness and sensitivity details

SI.2.1 Degree-preserving null model

The null comparison was designed to test whether the observed community structure was stronger than expected from degree configuration alone in the baseline country projections. For each analytical path, the observed country projection was rewired using a degree-preserving edge-swap procedure, after which the original edge weights were permuted and reassigned to the rewired topology. The null analysis was run for 40 replicates and is interpreted strictly as a stronger-than-null test of structural differentiation. It does not support causal claims about policy diffusion, transition speed, or project realization.

Figure S1 visualizes the null distributions for the count and capacity views, and Table S3 reports the corresponding summary statistics. In both cases, observed modularity is well above the null expectation. The capacity path shows an observed modularity of 0.684 against a null mean of 0.314 with a standard deviation of 0.010, yielding a z-score of 37.108. The count path shows an observed modularity of 0.618 against a null

Table S2: Country filter and bin assignment audit.

Country / edge case	Record count	Capacity (MW)	Passes threshold	Gap from 30	Note
Panel A. Country-threshold edge cases					
Mauritania	30	53,264.9	True	0	retained edge case
Puerto Rico	30	854.9	True	0	retained edge case
Tunisia	31	7,133.6	True	1	retained edge case
Kenya	31	1,011.7	True	1	retained edge case
Myanmar	31	690.5	True	1	retained edge case
Cambodia	32	1,096.0	True	2	retained edge case
Ecuador	28	1,404.4	False	-2	removed edge case
Madagascar	29	800.4	False	-1	removed edge case

Panel B. Bin assignment summary				
Capacity bin	**Rows**	**Capacity (MW)**	**Countries active**	**Boundary rule**
1–5	55,743	107,413.2	94	$1 \leq \text{MW} < 5$
5–50	28,040	478,419.5	94	$5 \leq \text{MW} < 50$
50–500	14,131	1,758,342.1	90	$50 \leq \text{MW} < 500$
500+	1,028	1,228,864.0	50	$\text{MW} \geq 500$

Threshold-edge cases and bin-assignment totals confirm that the retained sample was constructed deterministically.

mean of 0.293 with a standard deviation of 0.010, yielding a z-score of 33.311. No null replicate in either view reached the observed modularity.

These results support the interpretation that the identified communities reflect stronger-than-null structural differentiation in the retained country space. At the same time, the null comparison should be read only at the level of structural organization: it tests whether the observed grouping is stronger than a degree-preserving baseline, not the mechanisms that generate it.

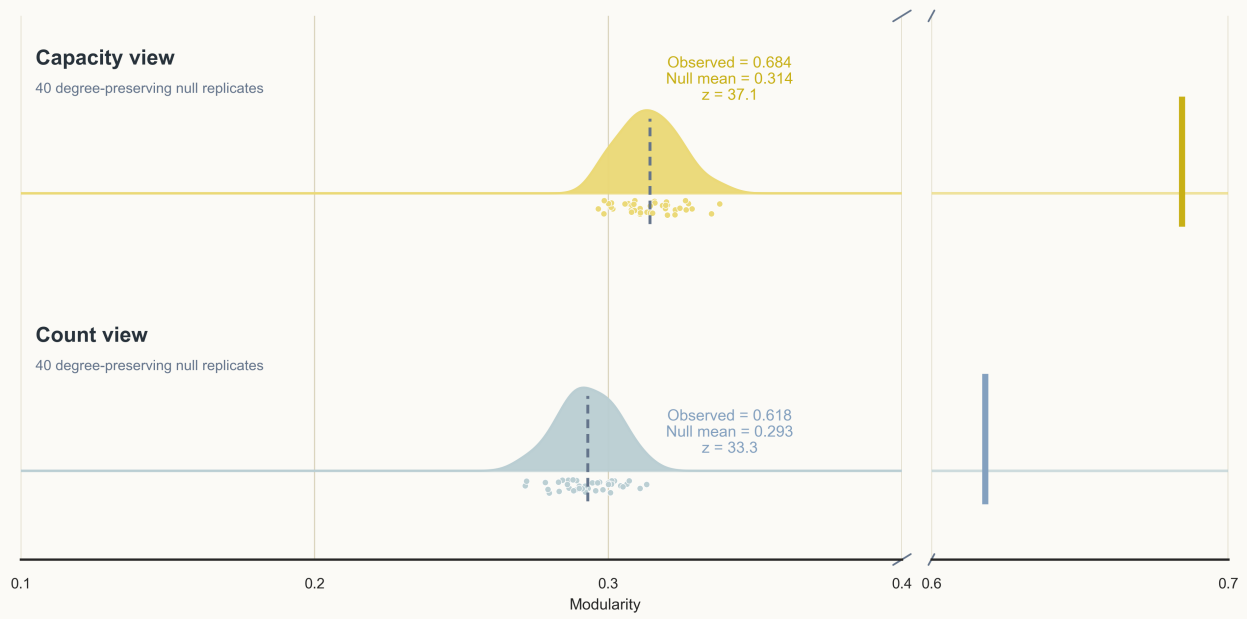

Figure S1: Null-model comparison. Observed modularity for the baseline capacity and count projections is well above the degree-preserving null distributions in both views, supporting a stronger-than-null interpretation of community structure.

Table S3: Null-model summary. Panel A reports the core null-comparison statistics for the baseline capacity and count projections. Panel B reports the null-distribution range summary. Panel C reports the corresponding null-run and graph configuration summary.

Panel A. Core null-comparison statistics

View	Observed modularity	Null mean	Null SD	z-score
Capacity	0.684	0.314	0.010	37.108
Count	0.618	0.293	0.010	33.311

Panel B. Null-distribution range summary

View	Null min	Null median	Null max	Null $\geq$ observed
Capacity	0.297	0.314	0.338	0/40
Count	0.272	0.292	0.313	0/40

Panel C. Null-run and graph summary

View	Null runs	Passes null gate	Nodes	Edges	Communities
Capacity	40	True	94	366	6
Count	40	True	94	405	6

The final column in Panel B reports the number of null replicates whose modularity was at least as large as the observed modularity. Because the null analysis used 40 replicates, this count is reported rather than an exact empirical p-value.

SI.2.2 Alternative binnings

Two alternative binning schemes were evaluated in addition to the baseline specification: a *low-end split* (1–5, 5–20, 20–100, 100+ MW) and a *giga split* (1–10, 10–100, 100–1000, 1000+ MW). Table S4 reports partition-level outputs for these alternative binnings and for the alternative country-inclusion rule, together with the corresponding top-country stability summaries.

Across both views, the alternative binnings retain the main operating, pipeline, and announcement-heavy structural axes at a broad level, although exact memberships shift. Under the giga split, the ARI relative to baseline is 0.517 for the capacity path and 0.505 for the count path. Under the low-end split, the ARI is 0.537 for the capacity path and 0.493 for the count path. These values indicate moderate reassignment at the level of exact membership, particularly where bin boundaries are redefined.

The main caution concerns the capacity path under the low-end split, where top-20 tier stability falls to 0.40. However, status stability for the top 20 countries remains 1.0, and the top-5 camp-signature overlap remains 0.80. This indicates that fine-grained boundary placement can shift when the lower end of the size distribution is subdivided more aggressively, while the higher-order structural axes remain visible. Thus, the sensitivity results should be interpreted as supporting continuity in broad structural axes rather than invariance in exact country membership.

SI.2.3 Country-filter sensitivity

A further descriptive sensitivity analysis replaced the baseline country filter with an alternative inclusion rule of total capacity $\geq$ 500 MW. This expanded the analytical base from 94 to 124 countries, with 99,338 rows and 3,658,030.6 MW represented under the wider filter. The purpose of this exercise was not to replace the baseline retained-scope rule, but to evaluate whether the main conclusions depended narrowly on the use of record count $\geq$ 30.

As shown in Table S4, the broader country universe retains the main structural logic in both analytical paths at the level of broad axes, although exact membership changes are again present. Relative to the baseline, the ARI remains 0.617 for the capacity path and 0.490 for the count path, while top-20 dominant-profile stability remains 1.0 in both views. Under the wider country set, the principal structural axes remain

Table S4: Sensitivity analysis under alternative specifications. Results are split into five compact panels for readability.

Panel 1. Alternative giga split

Scenario	View	Communities	Modularity	ARI vs. baseline
Giga split	Capacity	6	0.656	0.517
Giga split	Count	6	0.591	0.505

Panel 2. Alternative low-end split

Scenario	View	Communities	Modularity	ARI vs. baseline
Low-end split	Capacity	6	0.652	0.537
Low-end split	Count	7	0.568	0.493

Panel 3. Country-inclusion sensitivity

View	Countries	Rows	Capacity (MW)	Communities	Modularity	ARI vs. baseline
Capacity	124	99,338	3,658,030.6	8	0.730	0.617
Count	124	99,338	3,658,030.6	7	0.616	0.490

Panel 4. Top-country stability under alternative binning

Scenario	View	Status (top 20)	Tier (top 20)	Profile (top 20)	Overlap (top 5)
Giga split	Capacity	1.0	0.9	0.9	0.6
Giga split	Count	1.0	0.9	0.9	1.0
Low-end split	Capacity	1.0	0.4	0.4	0.8
Low-end split	Count	1.0	0.95	0.95	0.667

Panel 5. Top-country stability under country-inclusion sensitivity

Scenario	View	Status (top 20)	Tier (top 20)	Profile (top 20)	Overlap (top 5)
Total capacity $\geq$ 500 MW	Capacity	1.0	1.0	1.0	0.75
Total capacity $\geq$ 500 MW	Count	1.0	1.0	1.0	0.667

For the alternative binning schemes, the retained country set, row count, and total capacity remain unchanged relative to the baseline specification (94 countries, 98,942 rows, and 3,573,038.8 MW).
Status (top 20), Tier (top 20), Profile (top 20), and Overlap (top 5) summarize stability among leading countries or camp signatures. The table supports continuity in broad structural axes while allowing exact tier placement and country membership to shift across specifications.

recognizable: operating-heavy structure, mid-scale pipeline structure, and large-announcement structure. The baseline country rule is retained in the main paper because it is designed to stabilize country-level structural profiles rather than to privilege countries with the largest absolute capacity.

SI.3 Extended structural profiles and cross-view correspondence

SI.3.1 Extended dominant mode profiles for capacity camps

The main paper reports the dominant structural features of the six capacity camps. Table S5 provides the extended dominant-mode profiles used to support Figure 3 and Table 2 in the main text. For each camp, the table reports the number of countries, country share, graph-strength share, dominant mode, dominant-mode share, top-three modes, and combined top-three share.

These extended profiles show that the six capacity camps differ not only in their dominant stage and size layer, but also in the concentration and internal coherence of their dominant structures. For example, the 1–5 operating camp shows high internal concentration in operating modes, while the 500+ pre-construction / operating mixed camp combines very large projected capacity structure with a substantial system-weighted operating layer. The extended profiles therefore support the interpretation that the capacity camps describe distinct system-weighted configurations rather than a single monotonic scale gradient.

Table S5: Extended dominant-mode profiles for capacity camps. Panel A summarizes the size and weight distribution of the six camps. Panel B reports the dominant-mode profiles and top-three structural components of each camp.

Panel A. Capacity camp size and weight summary

Camp	Countries	Country share	Raw weight share	Graph-strength share
50–500 pre-construction	24	0.255	0.169	0.267
500+ pre-construction / operating mixed	18	0.191	0.642	0.157
50–500 announcement	17	0.181	0.040	0.175
1–5 operating	15	0.160	0.039	0.195
500+ announcement	13	0.138	0.098	0.134
5–50 operating	7	0.074	0.012	0.073

Panel B. Capacity camp dominant-mode profiles

Camp	Dominant mode	Top-3 modes	Top-3 share
50–500 pre-construction	pre-construction — 50–500	pre-construction — 50–500; pre-construction — 5–50; announced — 50–500	0.635
500+ pre-construction / operating mixed	pre-construction — 500+	pre-construction — 500+; operating — 50–500; pre-construction — 50–500	0.567
50–500 announcement	announced — 50–500	announced — 50–500; announced — 5–50; pre-construction — 50–500	0.673
1–5 operating	operating — 1–5	operating — 1–5; operating — 5–50; operating — 50–500	0.876
500+ announcement	announced — 500+	announced — 500+; pre-construction — 500+; announced — 50–500	0.708
5–50 operating	operating — 5–50	operating — 50–500; operating — 5–50; operating — 1–5	0.778

Extended profiles report dominant-mode concentration and top-three structural composition for each capacity camp.

SI.3.2 Extended dominant mode profiles for count ecologies

Table S6 provides the corresponding extended dominant-mode profiles for the five higher-order count ecologies used in the main text. These profiles show that the count path is anchored more strongly in small-operating modes than the capacity path, but that this small-operating profile space is internally heterogeneous. The

broad small-operating and high-purity small-operating ecologies, for example, share a common dominant mode but differ in profile purity and top-three concentration.

The count profiles also clarify that the higher-order ecology labels are not arbitrary. Each ecology is supported by a distinct dominant structure, and the blended or transition-oriented ecologies are not residual categories. Instead, they correspond to interpretable combinations of operating, pre-construction, and announced modes at specific size layers. Table S6 therefore provides extended numerical support for Figure 4.

Table S6: Extended dominant-mode profiles for count ecologies. Panel A summarizes the size and weight distribution of the five higher-order ecologies. Panel B reports the dominant-mode profiles and top-three structural components of each ecology.

Panel A. Count ecology size and weight summary

Ecology	**Source communities**	**Countries**	**Country share**	**Raw weight share**	**Graph-strength share**
Broad small-operating	1	25	0.266	0.468	0.265
Mixed small/mid-operating	2	23	0.245	0.109	0.250
High-purity small-operating	3, 6	22	0.234	0.324	0.267
Pipeline-transition	4	14	0.149	0.023	0.103
Small-operating mixed	5	10	0.106	0.076	0.115

Panel B. Count ecology dominant-mode profiles

Ecology	**Dominant mode**	**Top-3 modes**	**Top-3 share**
Broad small-operating	operating — 1–5	operating — 1–5; operating — 5–50; pre-construction — 50–500	0.669
Mixed small/mid-operating	operating — 5–50	operating — 5–50; operating — 1–5; operating — 50–500	0.760
High-purity small-operating	operating — 1–5	operating — 1–5; operating — 5–50; announced — 5–50	0.956
Pipeline-transition	pre-construction — 5–50	pre-construction — 5–50; announced — 50–500; pre-construction — 50–500	0.563
Small-operating mixed	operating — 1–5	operating — 1–5; operating — 5–50; operating — 50–500	0.900

Extended profiles report dominant-mode concentration and top-three structural composition for each higher-order count ecology.

SI.3.3 Cross-view correspondence support

The cross-view relationship between count ecologies and capacity camps is summarized in the main paper by the Sankey-style correspondence figure. Table S7 provides the corresponding numerical support in condensed form by listing the ranked flows from each count ecology into the capacity camps. The strongest one-to-one alignment is the mapping from High-purity small-operating ecology to the 1–5 operating camp. By contrast, Broad small-operating ecology and Mixed small/mid-operating ecology disperse across multiple pre-construction and announcement-heavy capacity camps.

This pattern supports the interpretation that count ecology and system-weighted structure are related but not equivalent. The ranked flow structure shows that some ecologies align tightly with one capacity camp, whereas others contain countries that share similar project-frequency environments but occupy different system-weighted positions. Table S7 therefore provides numerical support for the main paper's claim of partial alignment rather than one-to-one equivalence.

SI.3.4 Count merge crosswalk and audit

The count path originally yields six algorithmic communities. The main paper compresses these into five higher-order ecologies for narrative clarity, but does not overwrite the underlying partition. Table S8 documents this compression in three panels. Panel A provides the crosswalk from original count communities to the main-text ecologies. Panel B summarizes the similarity evidence for the strongest merge candidates. Panel C reports the resulting merge recommendations used to determine whether higher-order compression was justified.

Only one pair, communities 3 and 6, was strongly supported for merging. This pair shows an all-mode similarity of 0.983 and a discriminative-mode similarity of 0.984, with identical dominant mode and dominant status. All other plausible merge candidates were weaker and were therefore not used for main-text compression. The resulting 6-to-5 mapping should therefore be interpreted as an audit-supported narrative layer rather than a replacement of the algorithmic count partition.

Table S7: Cross-view correspondence between count ecologies and capacity camps.

Count ecology	Capacity camp	Country count	Rank within ecology
High-purity small-operating ecology	1–5 operating camp	14	1
High-purity small-operating ecology	50–500 pre-construction camp	4	2
High-purity small-operating ecology	500+ pre-construction / operating mixed camp	2	3
Broad small-operating ecology	50–500 pre-construction camp	8	1
Broad small-operating ecology	500+ pre-construction / operating mixed camp	7	2
Broad small-operating ecology	50–500 announcement camp	6	3
Broad small-operating ecology	500+ announcement camp	4	4
Small-operating mixed ecology	500+ pre-construction / operating mixed camp	3	1
Small-operating mixed ecology	50–500 announcement camp	3	1
Mixed small/mid-operating ecology	5–50 operating camp	6	1
Mixed small/mid-operating ecology	50–500 pre-construction camp	5	2
Mixed small/mid-operating ecology	500+ announcement camp	5	2
Pipeline-transition ecology	50–500 pre-construction camp	6	1
Pipeline-transition ecology	500+ pre-construction / operating mixed camp	3	2
Pipeline-transition ecology	50–500 announcement camp	3	2

Ranked flows list the leading capacity-camp destinations for each count ecology.

Table S8: Count-path community-to-ecology crosswalk and merge audit. Panel A provides the crosswalk from the original six count communities to the five main-text ecologies. Panel B summarizes similarity evidence for the strongest merge candidates. Panel C reports the resulting merge recommendations.

Panel A. Community-to-ecology crosswalk

Original community	Original draft label	Main-text ecology	Ecology label	Merge applied
1	Draft: operating 1–5 led	1	Broad small-operating ecology	False
2	Draft: operating 5–50 led	2	Mixed small/mid-operating ecology	False
3	Draft: operating 1–5 heavy	3	High-purity small-operating ecology	True
4	Draft: pre-construction / announced mixed (5–50-led)	4	Pipeline-transition ecology	False
5	Draft: operating 1–5 heavy	5	Small-operating mixed ecology	False
6	Draft: operating 1–5 heavy	3	High-purity small-operating ecology	True

Panel B. Merge similarity evidence

Left community	Right community	All-mode sim.	Disc.-mode sim.	Same dominant mode	Same dominant status
3	6	0.983	0.984	True	True
3	5	0.972	0.973	True	True
1	5	0.958	0.965	True	True
1	3	0.951	0.963	True	True
2	5	0.873	0.876	False	True
1	2	0.804	0.801	False	True

Panel C. Merge recommendation summary

Candidate pair	Support level	Recommendation	Note / rationale
3 + 6	recommended	merge for main-text compression	strongest all-mode and discriminative-mode similarity; same dominant mode and status
3 + 5	plausible	not used	high similarity, but weaker than 3 + 6 and not required for higher-order interpretation
1 + 5	plausible	not used	structurally related, but retained as separate ecologies for interpretive clarity
1 + 3	plausible	not used	both small-operating, but different purity justifies separation
2 + 5	not recommended	do not merge	weaker similarity and different dominant mode structure
1 + 2	not recommended	do not merge	captures substantively different small vs. mixed small/mid operating structure

Only communities 3 and 6 were used for higher-order compression in the main text; the original six-community partition is retained as the underlying algorithmic result.

SI.4 Country membership and transition transparency

This section provides country-level transparency for the baseline camp and ecology assignments and for the corresponding low-end binning sensitivity check. Tables S9 and S10 list the baseline assignments for all 94 retained countries in the capacity and count views. To keep the Supporting Information concise, these tables report only the country name and baseline assignment; the extended dominant-mode and weight profiles are reported separately in Tables S5 and S6. Tables S11 and S12 summarize country movements under the low-end alternative binning. They report transition counts, within-baseline-group shares, and the countries reassigned to a different dominant interpretation when applicable.

SI.4.1 Baseline country memberships

Table S9: Baseline capacity-camp membership for all retained countries.

Country	Baseline assignment
Albania	50–500 pre-construction
Algeria	50–500 pre-construction
Argentina	50–500 announcement
Armenia	50–500 announcement
Australia	500+ announcement
Austria	1–5 operating
Bangladesh	500+ pre-construction / operating mixed
Belgium	1–5 operating
Bosnia and Herzegovina	50–500 pre-construction
Brazil	50–500 pre-construction
Bulgaria	500+ pre-construction / operating mixed
Cambodia	50–500 announcement
Canada	50–500 pre-construction
Chile	50–500 pre-construction
China	500+ pre-construction / operating mixed
Colombia	50–500 pre-construction
Croatia	50–500 pre-construction
Cuba	500+ announcement
Cyprus	1–5 operating
Czech Republic	1–5 operating
Denmark	500+ pre-construction / operating mixed
Dominican Republic	50–500 pre-construction
Egypt	500+ announcement
El Salvador	5–50 operating
Estonia	1–5 operating
Finland	50–500 pre-construction
France	1–5 operating
Germany	1–5 operating
Ghana	500+ announcement
Greece	50–500 pre-construction
Honduras	50–500 announcement
Hungary	1–5 operating
India	500+ pre-construction / operating mixed
Indonesia	500+ announcement
Iran	500+ announcement
Iraq	500+ pre-construction / operating mixed
Ireland	50–500 announcement

Continued on next page

Table S9: Baseline capacity-camp membership for all retained countries. (continued)

Country	Baseline assignment
Israel	5–50 operating
Italy	50–500 pre-construction
Japan	1–5 operating
Jordan	500+ pre-construction / operating mixed
Kazakhstan	500+ announcement
Kenya	50–500 announcement
Latvia	50–500 announcement
Lithuania	50–500 pre-construction
Luxembourg	1–5 operating
Malaysia	500+ announcement
Mauritania	500+ announcement
Mexico	500+ pre-construction / operating mixed
Moldova	1–5 operating
Morocco	500+ pre-construction / operating mixed
Myanmar	5–50 operating
Namibia	500+ pre-construction / operating mixed
Nepal	50–500 pre-construction
Netherlands	5–50 operating
New Zealand	50–500 pre-construction
Nigeria	50–500 announcement
North Macedonia	500+ announcement
Norway	50–500 pre-construction
Oman	500+ pre-construction / operating mixed
Pakistan	500+ pre-construction / operating mixed
Panama	50–500 announcement
Peru	50–500 announcement
Philippines	500+ pre-construction / operating mixed
Poland	1–5 operating
Portugal	50–500 pre-construction
Puerto Rico	50–500 pre-construction
Romania	50–500 pre-construction
Russia	5–50 operating
Réunion	1–5 operating
Saudi Arabia	500+ pre-construction / operating mixed
Serbia	50–500 announcement
Singapore	50–500 pre-construction
Slovakia	1–5 operating
Slovenia	50–500 announcement
South Africa	500+ announcement
South Korea	1–5 operating
Spain	50–500 pre-construction
Sri Lanka	50–500 announcement
Sweden	50–500 pre-construction
Switzerland	50–500 pre-construction
Syria	50–500 announcement
Taiwan	500+ pre-construction / operating mixed
Thailand	50–500 announcement
Tunisia	500+ pre-construction / operating mixed
Türkiye	5–50 operating

Continued on next page

Table S9: Baseline capacity-camp membership for all retained countries. (continued)

Country	Baseline assignment
Ukraine	5–50 operating
United Arab Emirates	500+ pre-construction / operating mixed
United Kingdom	50–500 pre-construction
United States	500+ pre-construction / operating mixed
Uzbekistan	500+ announcement
Vietnam	50–500 announcement
Zambia	500+ announcement
Zimbabwe	50–500 announcement

SI.4.2 Baseline count-ecology memberships

Table S10: Baseline count-ecology membership for all retained countries.

Country	Baseline assignment
Albania	Broad small-operating
Algeria	Mixed small/mid-operating
Argentina	Mixed small/mid-operating
Armenia	Mixed small/mid-operating
Australia	Broad small-operating
Austria	High-purity small-operating
Bangladesh	Broad small-operating
Belgium	High-purity small-operating
Bosnia and Herzegovina	Broad small-operating
Brazil	Broad small-operating
Bulgaria	Small-operating mixed
Cambodia	Broad small-operating
Canada	Mixed small/mid-operating
Chile	Broad small-operating
China	Broad small-operating
Colombia	Pipeline-transition
Croatia	Broad small-operating
Cuba	Mixed small/mid-operating
Cyprus	High-purity small-operating
Czech Republic	High-purity small-operating
Denmark	Small-operating mixed
Dominican Republic	Pipeline-transition
Egypt	Mixed small/mid-operating
El Salvador	Mixed small/mid-operating
Estonia	High-purity small-operating
Finland	Pipeline-transition
France	Small-operating mixed
Germany	High-purity small-operating
Ghana	Pipeline-transition
Greece	Broad small-operating
Honduras	Mixed small/mid-operating
Hungary	High-purity small-operating
India	Mixed small/mid-operating
Indonesia	Small-operating mixed

Continued on next page

Table S10: Baseline count-ecology membership for all retained countries. (continued)

Country	Baseline assignment
Iran	Mixed small/mid-operating
Iraq	Pipeline-transition
Ireland	Pipeline-transition
Israel	Mixed small/mid-operating
Italy	High-purity small-operating
Japan	High-purity small-operating
Jordan	Small-operating mixed
Kazakhstan	Mixed small/mid-operating
Kenya	Small-operating mixed
Latvia	Small-operating mixed
Lithuania	High-purity small-operating
Luxembourg	High-purity small-operating
Malaysia	Mixed small/mid-operating
Mauritania	Broad small-operating
Mexico	Mixed small/mid-operating
Moldova	High-purity small-operating
Morocco	Broad small-operating
Myanmar	Mixed small/mid-operating
Namibia	Mixed small/mid-operating
Nepal	Pipeline-transition
Netherlands	Mixed small/mid-operating
New Zealand	Pipeline-transition
Nigeria	Broad small-operating
North Macedonia	High-purity small-operating
Norway	Pipeline-transition
Oman	Broad small-operating
Pakistan	High-purity small-operating
Panama	Mixed small/mid-operating
Peru	Pipeline-transition
Philippines	Pipeline-transition
Poland	High-purity small-operating
Portugal	Mixed small/mid-operating
Puerto Rico	Mixed small/mid-operating
Romania	Small-operating mixed
Russia	Mixed small/mid-operating
Réunion	High-purity small-operating
Saudi Arabia	Broad small-operating
Serbia	Broad small-operating
Singapore	High-purity small-operating
Slovakia	High-purity small-operating
Slovenia	High-purity small-operating
South Africa	Broad small-operating
South Korea	High-purity small-operating
Spain	Broad small-operating
Sri Lanka	Broad small-operating
Sweden	Broad small-operating
Switzerland	High-purity small-operating
Syria	Broad small-operating
Taiwan	High-purity small-operating

Continued on next page

Table S10: Baseline count-ecology membership for all retained countries. (continued)

Country	Baseline assignment
Thailand	Small-operating mixed
Tunisia	Pipeline-transition
Türkiye	Small-operating mixed
Ukraine	Mixed small/mid-operating
United Arab Emirates	Broad small-operating
United Kingdom	Mixed small/mid-operating
United States	Broad small-operating
Uzbekistan	Broad small-operating
Vietnam	Broad small-operating
Zambia	Pipeline-transition
Zimbabwe	Pipeline-transition

SI.4.3 Transition summaries under the low-end alternative binning

Tables S11 and S12 crosswalk the baseline assignments to the corresponding groups under the low-end alternative binning specification. Stable transitions are summarized by counts and shares; reassigned countries are listed explicitly to make sensitivity at the country level directly inspectable without reproducing full alternative membership tables.

Table S11: Capacity-view transition summary under the low-end alternative binning specification.

Baseline group	Alternative group	Countries	Share	Reassigned countries, if any
50–500 pre-construction	Alt-Group 3 (pre-construction 20–100; n=16)	16	0.67	—
50–500 pre-construction	Alt-Group 2 (pre-construction 100+; n=19)	6	0.25	Algeria, Canada, Chile, Lithuania, Portugal, Romania
50–500 pre-construction	Alt-Group 1 (announced 100+; n=27)	1	0.04	Nepal
50–500 pre-construction	Alt-Group 6 (announced 100+; n=7)	1	0.04	New Zealand
500+ pre-construction / operating mixed	Alt-Group 2 (pre-construction 100+; n=19)	12	0.67	—
500+ pre-construction / operating mixed	Alt-Group 5 (operating 20–100; n=9)	4	0.22	Bulgaria, Denmark, Jordan, Taiwan
500+ pre-construction / operating mixed	Alt-Group 6 (announced 100+; n=7)	2	0.11	Bangladesh, Philippines
50–500 announcement	Alt-Group 1 (announced 100+; n=27)	17	1.00	—
1–5 operating	Alt-Group 4 (operating 1–5; n=16)	15	1.00	—
500+ announcement	Alt-Group 1 (announced 100+; n=27)	8	0.62	—
500+ announcement	Alt-Group 6 (announced 100+; n=7)	4	0.31	Australia, Egypt, South Africa, Zambia
500+ announcement	Alt-Group 2 (pre-construction 100+; n=19)	1	0.08	Uzbekistan
5–50 operating	Alt-Group 5 (operating 20–100; n=9)	5	0.71	—

Continued on next page

Table S11: Capacity-view transition summary under the low-end alternative binning specification. (continued)

Baseline group	Alternative group	Countries	Share	Reassigned countries, if any
5–50 operating	Alt-Group 1 (announced 100+; n=27)	1	0.14	Myanmar
5–50 operating	Alt-Group 4 (operating 1–5; n=16)	1	0.14	Türkiye

Table S12: Count-view transition summary under the low-end alternative binning specification.

Baseline group	Alternative group	Countries	Share	Reassigned countries, if any
Broad small-operating	Alt-Group 2 (operating 1–5; n=17)	8	0.32	—
Broad small-operating	Alt-Group 7 (operating 1–5; n=8)	7	0.28	Albania, Bosnia and Herzegovina, Brazil, Croatia, Serbia, Spain, Sweden
Broad small-operating	Alt-Group 1 (pre-construction 100+; n=19)	4	0.16	Bangladesh, Morocco, Oman, Saudi Arabia
Broad small-operating	Alt-Group 5 (operating 1–5; n=10)	2	0.08	Cambodia, China
Broad small-operating	Alt-Group 6 (operating 1–5; n=10)	2	0.08	Greece, Sri Lanka
Broad small-operating	Alt-Group 3 (operating 5-20; n=15)	1	0.04	Chile
Broad small-operating	Alt-Group 4 (operating 1–5; n=15)	1	0.04	United States
Mixed small/mid-operating	Alt-Group 3 (operating 5-20; n=15)	14	0.61	—
Mixed small/mid-operating	Alt-Group 5 (operating 1–5; n=10)	7	0.30	Cuba, Egypt, Honduras, Kazakhstan, Malaysia, Mexico, Myanmar
Mixed small/mid-operating	Alt-Group 1 (pre-construction 100+; n=19)	1	0.04	Argentina
Mixed small/mid-operating	Alt-Group 7 (operating 1–5; n=8)	1	0.04	Puerto Rico
High-purity small-operating	Alt-Group 4 (operating 1–5; n=15)	14	0.64	—
High-purity small-operating	Alt-Group 6 (operating 1–5; n=10)	8	0.36	Austria, Japan, Luxembourg, Poland, Singapore, Slovenia, South Korea, Switzerland
Pipeline-transition	Alt-Group 1 (pre-construction 100+; n=19)	14	1.00	—
Small-operating mixed	Alt-Group 2 (operating 1–5; n=17)	9	0.90	—
Small-operating mixed	Alt-Group 5 (operating 1–5; n=10)	1	0.10	Kenya

SI.5 Projection-parameter sensitivity

To evaluate whether the baseline groupings depend narrowly on the selected projection sparsity or Louvain resolution setting, an additional parameter sensitivity check was conducted for both the capacity and count

views. The baseline specification uses top-$k = 6$ country–country ties per node and Louvain resolution $= 1.0$. Alternative projection sparsity settings were evaluated at top-$k = 5$, top-$k = 7$, and top-$k = 8$, while alternative resolution settings were evaluated at resolution $= 0.8$ and resolution $= 1.2$.

Table S13 summarizes the results. The capacity view shows stronger agreement with the baseline than the count view across the tested variants, with ARI values ranging from 0.844 to 0.891 outside the baseline. The count view is more sensitive in exact membership, especially as the projection becomes denser, with ARI values ranging from 0.578 to 0.725 outside the resolution $= 0.8$ case, which reproduces the baseline count partition. However, top-20 status and tier stability remain 1.0 across all tested variants, and the main operating, pipeline, and announcement-heavy structural axes remain identifiable in both views. These results indicate that projection-parameter changes affect exact boundaries more than the higher-order structural differentiation reported in the main text.

Table S13: Projection-parameter sensitivity under alternative top-k and Louvain resolution settings.

View	**Variant**	**ARI**	**Status/Tier**	**Overlap**	**Main structure**
Capacity	Top-$k = 5$	0.860	1.00/1.00	1.00	Identifiable
Capacity	Top-$k = 7$	0.866	1.00/1.00	1.00	Identifiable
Capacity	Top-$k = 8$	0.844	1.00/1.00	0.80	Identifiable
Capacity	Resolution $= 0.8$	0.891	1.00/1.00	0.80	Identifiable
Capacity	Resolution $= 1.2$	0.884	1.00/1.00	1.00	Identifiable
Count	Top-$k = 5$	0.618	1.00/1.00	1.00	Identifiable
Count	Top-$k = 7$	0.678	1.00/1.00	0.67	Identifiable
Count	Top-$k = 8$	0.578	1.00/1.00	0.67	Identifiable
Count	Resolution $= 0.8$	1.000	1.00/1.00	1.00	Identifiable
Count	Resolution $= 1.2$	0.725	1.00/1.00	1.00	Identifiable

ARI = adjusted Rand index relative to the baseline partition within the same view. Status/Tier reports top-20 status stability and top-20 tier stability, respectively. Overlap reports top-5 camp-signature overlap. "Identifiable" indicates that the main operating, pipeline, and announcement-heavy structural axes remain recognizable under the tested variant; it should not be interpreted as exact country-membership invariance.