Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Two Projections of the Belt and Road: Mapping to Itself and to the World

Fanze Meng

uclzfme@ucl.ac.uk

University College London https://orcid.org/0009-0000-0488-581X

Research Article

Keywords: Belt and Road Initiative, China, regionalism, spatial analysis, strategic communication, political economy of infrastructure

Posted Date: November 11th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11686494/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License. Read Full License

Additional Declarations: The authors declare no competing interests.

Abstract

Beijing maintains two versions of the Belt and Road, one for audiences inside China and one for audiences beyond it. Whether the two differ in a structured way is the question answered here. The study recovers the boundary of Belt and Road space from the BRI Portal ecosystem by separating the internal, Chinese-language projection from the outward projection in English, Arabic, French, Russian and Spanish, with the source architecture held constant and only the audience varied. Country prominence is measured with a composite index over project counts, sectoral breadth, hard-connectivity share and path dependence, tested by scaling to GDP, re-weighting, a coverage-matched null and portal decomposition, and validated against AidData's record of Chinese official finance. The two projections correlate at Spearman ρ = 0.61 and do not coincide. Outward materials give Southeast Asian operational weight over to North African and CIS visibility, amplify Egypt and the Gulf, and damp Indonesia, the most deeply embedded host in the internal corpus, in all five outward languages. Portal prominence still tracks the financial scale of documented engagement (ρ = 0.60 across 125 mutually confirmed hosts, about 0.43 net of economy size). Beijing's official cartography of the Belt and Road is therefore two maps with a structured divergence between them. The maps give the regionalism literature's graded conception of regional formation an operational footing, supply a replicable boundary of officially represented BRI space in place of inherited country lists, and carry a dual-projection method portable to any regional project whose membership is politically constructed.

1. Introduction

Where does the Belt and Road begin and end, and does Beijing answer that question the same way to itself as it does to the world? The first half is familiar. The second is what this article pursues. The Belt and Road Initiative (BRI) was never constituted as a treaty organization with a charter, a membership roster, or a threshold for accession. It emerged instead as a flexible political project articulated through instruments such as memoranda, project announcements, financing arrangements, sectoral branding and a sprawling official communication apparatus (Cai 2017). Its boundaries are not missing data about a stable institution; they are internal to the object, and their openness has been politically useful. The initiative scales outward without legal commitment and is recalibrated as priorities and audiences shift (Nedopil Wang 2023; Li and Wang 2025). This adaptive recalibration is not peculiar to the BRI; it reflects the campaign-style, experiment-and-adjust repertoire long characteristic of Chinese policy implementation (Heilmann and Perry 2011).

That openness has left the field inheriting its operational geography without ever defending it. Counts of “BRI countries” diverge because they answer different questions: has a state signed a document, hosted a project, received finance later labeled BRI by outside observers, or been routinely named by Beijing as part of the initiative (World Bank 2019; Malik et al. 2021)? These are overlapping but non-identical universes, and conflating them is costly: a signatory list captures diplomatic embrace without operational depth, a finance dataset captures activity without telling us whether Beijing frames it as BRI, and a country basket assembled for one purpose is silently reused for another. The boundary is equally slippery in the field that studies the initiative routinely: work published in this journal shows how readily it absorbs projects conceived elsewhere, with a regional East African corridor initiated by national governments now carried as a BRI “project catalysing connectivity” (Chiyemura, Gambino and Zajontz 2023, p. 108). Work in this journal has taken the opposite path, widening the measure rather than fixing the boundary: engagement with the initiative is scored, in that work, on three elements, memorandum signature, participation in Belt and Road forums and membership of the Asian Infrastructure Investment Bank (Sims, Lee and Lee 2025, p. 9). The wider measure buys breadth at the cost of a stated boundary, which is the cost this article sets

out to avoid. As Hillman puts it, “it is impossible to fathom the BRI without a map” (Hillman 2020, p. 4), yet the initiative is “opaque and vague by design” (Hillman 2020, p. 7), so the map has to be built. The regionalism literature explains why this is not a trivial anomaly. Regions are politically constructed spaces (Mansfield and Milner 1999; Katzenstein 2005), and Hettne and Söderbaum’s (2000) concept of regionness treats regional formation as a matter of degree. A parallel sociology of European integration makes the same move for a treaty-based region, treating its boundaries as socially produced and functionally porous (Favell and Guiraudon 2009; Fligstein 2008), an insight since formalised in work on differentiated integration (Schimmelfennig and Winzen 2020). The contrast with treaty-based regions is therefore less sharp than it appears: a method for recovering politically-constructed boundaries has purchase well beyond an initiative as informal as the BRI. This article maps one specific space, how Beijing officially projects BRI space, which is empirically tractable. What the BRI “really is” remains an ontological question the design does not settle. The official BRI Portal ecosystem is the state’s own curated record of the initiative, and it can be read in two registers. Domestic-facing Chinese-language materials reconstruct the project universe Beijing records for itself; multilingual outward-facing materials reveal what it chooses to foreground for foreign audiences. Mapped separately and then compared, they pose a question no single source can: whether the official cartography of the BRI is one map or two.

It is two. Holding the Portal architecture constant and varying only the audience, the internal and outward projections correlate but diverge in structured ways ($\rho = 0.61$). The divergence is not noise around a shared mean. Southeast Asian regional weight is traded outward for North African and CIS visibility, Egypt and the Gulf are amplified, and Indonesia, first by a wide margin in the internal corpus, falls out of the top tier in every one of the five outward language portals. Three checks establish that this is real, not mechanical. It survives language-by-language decomposition, so it is not a by-product of each portal serving its natural regional audience. It survives a coverage-matched null and alternative index weightings, so it is not an artifact of construction. And it survives scaling to economic size. The damping is not proportionality to a small economy. Read against an independent benchmark, AidData’s reconstruction of Chinese official finance, portal prominence moves with the financial scale of documented engagement ($\rho = 0.60$ raw across 125 mutually-confirmed hosts, about 0.43 net of economy size), so the maps are anchored in material engagement and not in communicative selection alone. In the literature on adaptive Chinese governance, divergence of this kind is what one would expect. Ambiguity and clarity are combined adaptively in Chinese policy practice (Ang 2023); BRI narratives are elastic and audience-sensitive (Zeng 2017, 2019; Yang and Van Gorp 2023; Cheng and Zeng 2023). Differentiated outward projection is, on this reading, audience-differentiated projection to be measured rather than inconsistency to be corrected, a candidate instance of the image-management authoritarian states conduct abroad (Dukalskis 2021; Repnikova 2022). Prominence and projection therefore bear on influence and also on the closeness and relative importance of a relationship as Beijing construes it; because rhetorical and practical closeness need not coincide, the gap between them is itself analytically informative.

The article’s central contribution is descriptive: it establishes that the structured divergence exists and is not reducible to audience specialization, index construction, or economic size. How to explain it (deliberate audience management, reputational filtering, temporal asymmetry between the corpora, or the uneven availability of language-specific materials) is a question the design surfaces and narrows but does not adjudicate (Section 5); the cross-linguistic consistency of the damping and amplification (Section 4.5) is the single result that resists a pure production-artifact reading. Recovering it required two supporting results, which are contributions in their own right. The first is a defensible de facto boundary of officially represented BRI space: treating the Portal as the authoritative record of state-confirmed BRI space yields an explicit, replicable operational scope that replaces the inherited baskets circulating in the field, and that future research can adopt, contest or refine and can re-run to track how

official BRI space shifts over time. The second is the method behind it: primary-source architecture held constant, audience varied, the resulting maps validated against independent evidence. That method generalizes to any regional project that lacks treaty-fixed membership but maintains an official representational architecture. The wider purpose is to reorient BRI influence analysis from counts of membership toward differentiated prominence, projection, and externally-confirmed material engagement.

2. Literature Review

2.1 Ways of mapping the BRI

Across the field, most attempts to map the BRI operationalise it through external indicators, and they answer a different question from the one posed here. The largest family is finance-led. AidData's reconstruction of Chinese development finance traces where state-linked capital has gone (Malik et al. 2021; Burgess and Custer 2022), while World Bank and related modeling estimates the trade-cost and connectivity implications of a presumed BRI corridor network (de Soyres et al. 2018; World Bank 2019; Maliszewska and van der Mensbrugghe 2019; Haralambides and Merk 2020). A second family documents concrete spatial architectures such as railways, ports, financing flows and procurement (Liu, Zhang and Xiong 2020; Dollar 2019; Hurley, Morris and Portelance 2018). A third operationalises the BRI through agreements or participation labels (Wang 2021; Russel and Berger 2019), and a fourth proceeds case by case: Central and Eastern Europe (Vangeli 2020a, 2020b), Southeast Asia (Bharti and Kumari 2024), Africa (Dollar 2019; Yan and Sautman 2023). Each is indispensable for tracing material exposure, but none recovers Beijing's own selective official cartography: a country may loom large in a finance database yet be muted in Beijing's BRI communication, and vice versa.

The article draws on one of these, AidData's finance record, but not as its source for mapping BRI space. AidData enters later (Section 4) as an independent yardstick against which to test the portal-derived maps; why it is preferred to the finance-flow and agreement-based alternatives is set out with the methodology (Section 3). What the field lacks is a method that recovers official self-representation at the project level and compares it across audiences, not more data. Although BRI discourse has been studied at length, the official portal ecosystem has not, to our knowledge, been used as a project-level source for mapping BRI space: existing portal-facing work analyses its rhetoric and has not built a spatial corpus from it (Yang and Van Gorp 2023; Cheng and Zeng 2023; Zehra et al. 2025).

2.2 The BRI as a region-making project

In the regionalism literature the BRI's missing boundary is less an anomaly than an expectation. Regions are politically constructed formations (Mansfield and Milner 1999); Hettne and Söderbaum's (2000) concept of regionness treats them as graded and multidimensional, able to deepen, stall or fragment (Söderbaum 2016). On this view the BRI is best read as a region-making project rather than as a region with an uncertain edge, an uneven, selective construction of space whose coherence varies across places and sectors (Vangeli 2020a).

That reframing carries a direct methodological consequence. If the initiative orders space by degree, a binary map of who is "in" cannot carry an influence analysis; what matters is how intensely, broadly and durably particular countries are drawn into the initiative's official architecture. This is what motivates moving from presence to prominence. Once prominence is the object, it also motivates moving from one map to two, since a region-making project narrated to different audiences need not be made the same way for each. In that schema the BRI sits low on the ladder, a regional complex of overlapping, China-centered ties, and what the Prominence Index measures is the

representational side of this region-making: how intensely the official cartography narrates a country's belonging, not the achieved density of regional integration.

2.3 Why official self-representation is worth mapping

If the preceding sections show that existing maps inherit their country sets from outside Beijing's own representation, and that the initiative is best read as graded region-making, a further literature explains why that representation is itself worth mapping: official visibility is a governed, strategic act. The BRI's elasticity has been politically useful, functioning as slogan, umbrella, signalling device and coordination platform (Cai 2017; He 2019; Rolland 2022), and its narratives are broad, flexible and audience-sensitive (Zeng 2017, 2019; Yang and Van Gorp 2023; Cheng and Zeng 2023). Yilmaz and Li (2020, p. 395) go further in the other direction, arguing that BRI-led globalization may dampen and dilute China's participation in East Asian regionalization. Ang (2023) sharpens the point: ambiguity and clarity in Chinese policy communication combine adaptively, so differentiated communication across audiences should be expected, not read as incoherence.

The governance literature points the same way. Chinese SOEs remain embedded in Party-state hierarchies and serve strategic as well as commercial functions (Leutert 2020), so a portal corpus that privileges state-confirmed, SOE-centered projects captures a politically meaningful subset of BRI space. Large-scale Chinese overseas strategy is implemented through a fragmented but politically directed architecture, mobilised across provinces, ministries and state firms (Hameiri and Jones 2015; Jones and Hameiri 2021; Ye 2020). Two implications follow. Official portal materials are not transparent mirrors of an underlying project universe; they are curated records shaped by the communicative priorities of the state and its linked actors. And divergence across audience-facing materials should not be pre-judged as inconsistency. Selective disclosure may itself be how governance is exercised. This cuts two ways for the outward register in particular: it is shaped by a coordinated external-communication apparatus that would make audience differentiation deliberate (Brady 2015), yet it is also aggregated from many semi-autonomous producers whose uncoordinated output could generate divergence as a byproduct. The article therefore treats divergence as a possibility the literature generates and the design tests.

The gap is now precise. Work on datasets and project tracking shows how to operationalise Chinese activity through external indicators; work on ambiguity and governance shows why official communication is likely to be selective and audience-sensitive. What is missing is a method that integrates the two by comparing Beijing's domestic-facing and outward-facing projections within a single primary-source architecture, which is what the dual-projection design supplies.

3. Methodology

The design measures one thing and brackets two others, and the three are worth separating before the source and the index are introduced. Prominence is how heavily Beijing foregrounds a given country in its official representation of BRI space, the quantity the Prominence Index operationalises (Section 3.3). It is a property of representation, not of underlying activity: a country's prominence can rise or fall while its material footprint holds constant. Projection is prominence rendered for a particular audience; the internal and outward projections are two renderings of the same represented space, so holding the source architecture constant and varying only the audience isolates the representational act itself.

Influence, the design does not measure. Prominence bears on the closeness of a relationship as Beijing construes it, what it wishes foregrounded as central, but closeness and prominence can come apart, as the damped sensitive partners and amplified showcase states of Section 4 make plain; and neither translates directly into leverage over a

host, which would require evidence on reception, project performance and bargaining that representation alone cannot supply (Section 5). The point of holding the source constant is to measure the first term cleanly and to bracket the last.

3.1 The portal as primary source

The article's corpus is the official BRI Portal ecosystem (yidaiyilu.gov.cn), treated as authoritative because it is the institutional architecture through which the Chinese state constructs and projects BRI space. Its masthead makes the chain explicit: supervised by the Office of the Leading Group for Promoting the Belt and Road Initiative, guided by the National Development and Reform Commission and Xinhua, and hosted by the China Economic Information Service and the State Information Center. Every link is a state agency, and these bodies sit at the center of China's policy-coordination and external-communication apparatus (Brady 2008). A project's appearance on the portal is therefore itself an act of state recognition: whether or not its text uses the phrase □□□□, the state-supervised portal has chosen to present it as BRI space. This logic differs from finance- or investment-led datasets, which reconstruct underlying activity whether or not Beijing highlights it; here, official curation is the phenomenon under study. The portal is authoritative for the question of official self-representation and bounded relative to the total universe of Chinese activity abroad. Inclusion does not exhaust a country's Chinese footprint, nor does exclusion prove its absence. The portal publishes no consolidated membership roster, country list, or map; the country-level corpus is compiled bottom-up from its project-reporting and news streams, which is why a disciplined consolidation pipeline is needed.

3.2 Two projections, one architecture

The corpus is read in two registers. The internal projection is built from domestic-facing Chinese-language materials: weekly project reports, the portal news stream, and a "small and beautiful" (□□□) thematic stream, the portal's post-2021 label for smaller livelihood projects. The outward projection is built from the multilingual portal in English, Arabic, French, Russian and Spanish. Holding the architecture constant and varying only the audience is what keeps any later divergence from being an artifact of comparing unlike sources. The five outward-language corpora were read and coded in the original language by native-speaking translators, not by machine translation; where the same project appeared in more than one language, entries were consolidated on shared portal project identifiers, with translated-title matching as a fallback. The outward materials are analyzed both in aggregate and stratified by language (Section 4.5). In total the corpus comprises 6,192 article-level observations, of which 4,216 are domestic-facing (3,697 weekly project-report entries, 316 news items, 203 small-and-beautiful items) and 1,976 span the five outward languages, which the consolidation pipeline reduces to 3,892 internal and 1,688 outward canonical projects. These materials accumulate press releases and reports spanning roughly 2015–2026, concentrated in 2018–2025; the corpus is thus not a single snapshot, a point that bears on the temporal comparison in Section 4.7.

Table 1
Corpus assembly: article-level observations and canonical projects.

Layer	Article-level observations	Canonical projects
Internal (domestic-facing)	4,216	3,892
Outward (five language portals)	1,976	1,688
Total	6,192	5,580

3.3 From articles to canonical projects, and the Prominence Index

Because a single project may be reported many times and across several languages, raw article counts would inflate the prominence of frequently-reported countries. A six-step consolidation pipeline drawn from the broader project design (URL normalisation, cross-language aggregation, project-name de-duplication, country-and-contractor verification, manual adjudication of ambiguous cases, and unique project-ID assignment) converts article-level observations into canonical project records using standard record-linkage techniques (Christen 2012). Binary presence then yields Map 0, the minimum operational baseline: which countries host at least one canonical project.

At the level of its rules, the consolidation is reproducible. Repeated weekly-report entries for one project are collapsed by fuzzy-matching normalised project names at a Levenshtein distance of 0.15, a threshold calibrated on a manually verified pilot of 200 matched pairs, and a candidate merge is confirmed only when host country, contractor, and project category also agree, the most complete record being retained as canonical; borderline matches (distance 0.15 to 0.25) are flagged to a disambiguation log and adjudicated by hand. The eight-variable coding of the outward corpus was conducted against a shared codebook by coders with native or near-native proficiency; a 10 percent double-coded sample returned Cohen's κ between 0.81 and 0.94, discrepancies resolved by a third coder. The de-duplicated project database, per-country BPI components, and disambiguation log are archived for replication. The 0.15 threshold is calibrated on the pilot and has not been validated against a held-out gold sample, so its linkage accuracy is a tuning estimate.

To move from presence to intensity, the article proposes the BRI Prominence Index (BPI), a country-level composite of four components, each min-max scaled to the unit interval (Nardo et al. 2008) and given equal weight: project count, sector diversity across coded project categories, hard-connectivity share and path-dependence. Equal weighting is a deliberate default rather than a claim about the components' relative importance, and Section 4.6 reports two alternatives that shift weight toward count or toward breadth. The four register, respectively, the frequency of official inclusion, the breadth of a country's portfolio, its tilt toward physical infrastructure, and its depth of lock-in. Path-dependence is an ordinal code (high, medium, low) assigned during manual inspection, capturing the degree to which a country's projects entail Chinese technical standards, software systems, or financing arrangements that raise a partner's switching costs, the lock-in at issue in the debt-sustainability debate, whose 'debt-trap' framing has itself been substantially qualified (Brautigam 2020). The hard-connectivity component records how Beijing's own representation skews toward physical infrastructure. It measures the official portrait, not an endorsement of the view that hard infrastructure matters most, and, being near-uniformly high across the corpus, it adds little to the ranking in practice. Each record is coded hard or soft on the portal's own □□□/□□□ designation where one is present, and on the sectoral default otherwise; the explicit designation therefore overrides the sector, which is why the hard share exceeds what a sector-by-sector reading of Table 5 yields. Whether prominence tracks financial commitment is tested separately, by scaling against host-economy size (Section 4.4) and against AidData's finance record (Section 4.7); the index itself carries no financial component. Because the components are largely collinear, the rankings are robust to the composite's design, as the reweighting checks and component diagnostics in Section 4.6 confirm.

3.4 Validation strategy and inference

Two procedures guard the findings. First, because an observed divergence between the projections could in principle arise from sampling alone, a coverage-matched null asks how strongly the two projections would correlate

if they were merely different-sized samples of one underlying country prominence. Each country's pooled internal-and-outward total defines that latent prominence, from which an internal and an outward sample are drawn at the two corpora's sizes and the cross-projection rank correlation recomputed across 3,000 simulations (Good 2005). Second, to test whether portal prominence corresponds to material engagement or only to communicative selection, the maps are cross-validated against AidData's BRI-Tagged Chinese Official Finance Dataset (Burgess and Custer 2022). AidData, a development-finance research lab at William & Mary that reconstructs Chinese state financing independently of Beijing, is the appropriate benchmark on four counts: it is epistemically independent of the portal, project-level, explicitly BRI-coded against the five thematic priorities of Xi's 2013 BRI speeches, and methodologically replicable. The finance-flow and agreement-based alternatives each fail at least one test: money-flow data is not BRI-specific, and memorandum lists record signatures and nothing at project level.

4. Analysis

4.1 The outer boundary (Map 0)

At the binary level the portal corpus covers 158 host countries (Fig. 1), broader than the memorandum-of-understanding counts cited in policy work, but narrower than the set of all states with significant Chinese commercial presence. Consistent with Fig. 1, it excludes most of Western Europe, North America, Japan, South Korea and India, which enter the corpus, if at all, only as termini of cross-border infrastructure. Of the 158, 133 appear in both projections, 18 only internally, and 7 only outwardly, and the asymmetry is a first signal that the two projections record different things. Internal-only states include small Pacific islands and Western rail-termini; the outward-only set (North Korea, Lebanon, Palestine, Eritrea, and several small Caribbean and European states) appears for diplomatic or symbolic reasons, with no recorded project activity. Map 0 is a useful outer envelope, but presence alone cannot carry an influence analysis; for that we turn to prominence.

Table 2
Coverage of the portal corpus by layer.

Layer	Host countries	Canonical projects
Both projections	133	3,892 internal; 1,688 outward
Internal only	18	—
Outward only	7	—
Union (Map 0)	158	—

4.2 The internal projection (Map A)

Map A (Fig. 2) reconstructs Beijing's internal projection from the domestic-facing corpus. It is structurally infrastructural and Asia-centered: 92.0 percent of the 3,892 canonical internal projects are hard-connectivity (transport 32.0 percent, energy 30.0, civil construction 16.2, industrial 7.4, water 5.2, mining 1.8), with softer or symbolic categories together below 6 percent. Hard-connectivity share contributes little to the country ranking (Section 4.6), so the index does not merely re-describe a hard/soft preference.

By a wide margin Indonesia leads (BPI = 0.92, 207 projects, ten sectors), ahead of a tier combining large counts with broad sectoral spread: Malaysia, Nigeria, Tanzania, Cambodia, Uzbekistan, Laos, Bangladesh, Kenya, Pakistan, Serbia, Saudi Arabia and Ethiopia (0.73–0.80). While concentration is real, it is not extreme: the top ten internal countries account for 32.2 percent of canonical projects, leaving a populated long tail. That breadth is itself

suggestive: a heavily concentrated footprint would point toward a small set of geopolitically-selected partners, whereas a populated long tail is more consistent with broad, economically motivated engagement. Adjudicating motive needs the finance- and motive-level evidence discussed in Section 5.

Table 3
The top twenty host countries in the internal projection (Map A), with the two alternative weightings.

			Corpus		Prominence under re-weighting		
Rank	**Country**	**BPI**	**Projects**	**Sectors**	**Region**	**BPI count-heavy**	**BPI breadth-heavy**
1	IDN	0.92	207	9	Southeast Asia	0.94	0.91
2	MYS	0.80	124	10	Southeast Asia	0.76	0.81
3	NGA	0.80	111	9	Sub-Saharan Africa	0.74	0.81
4	TZA	0.79	118	10	Sub-Saharan Africa	0.74	0.81
5	KHM	0.79	115	11	Southeast Asia	0.74	0.81
6	UZB	0.79	99	10	Central Asia	0.72	0.81
7	LAO	0.78	114	11	Southeast Asia	0.74	0.80
8	BGD	0.77	112	9	South Asia	0.73	0.79
9	KEN	0.77	105	10	Sub-Saharan Africa	0.72	0.79
10	PAK	0.74	107	9	South Asia	0.70	0.76
11	SRB	0.74	78	9	Central-Eastern Europe and Balkans	0.67	0.76
12	SAU	0.74	136	6	Middle East	0.72	0.73
13	ETH	0.73	88	11	Sub-Saharan Africa	0.67	0.76
14	GIN	0.71	44	9	Sub-Saharan Africa	0.61	0.74
15	PHL	0.70	97	6	Southeast Asia	0.66	0.70
16	AGO	0.68	69	8	Sub-Saharan Africa	0.61	0.70
17	THA	0.68	71	8	Southeast Asia	0.61	0.69
18	KAZ	0.67	57	6	Central Asia	0.59	0.68
19	ZMB	0.67	35	8	Sub-Saharan Africa	0.57	0.69
20	ARE	0.67	76	7	Middle East	0.61	0.68
Host countries			151				
Canonical projects			3,892				

Note. BPI is the equal-weighted mean of four min-max scaled components; the two right-hand columns re-weight them.

The internal projection’s regional order is led by Sub-Saharan Africa (30.8 percent of projects across forty-six hosts) and Southeast Asia (22.9 percent across eleven), followed by the Middle East (9.9), Latin America (8.5), South Asia (8.4) and Central Asia (6.6); Central-Eastern Europe, North Africa, Oceania and the CIS form a thin third tier, while Western Europe, the Caucasus, East Asia and North America register only at the margin. Although Sub-Saharan Africa leads by project count, its mean BPI (0.51) is a wide spread of small footprints, whereas the Southeast Asian mean (0.69) stands above every other region, the clearest single-region signal in the internal map, and the one the outward projection most sharply revises.

4.3 The outward projection (Map B)

In Map B (Fig. 3) the picture reorders. Egypt now leads (BPI = 0.92), followed by Pakistan and Kenya (0.83), then Ethiopia, Cambodia, Laos, Saudi Arabia, Serbia, Russia and Tanzania. Indonesia falls to 0.60; Malaysia leaves the top twenty entirely; Egypt, Brazil, Argentina, Peru, South Africa, Algeria, Russia and Sri Lanka enter it. The top-twenty overlap between the projections is only 11 of 20.

Table 4
The top twenty host countries in the outward projection (Map B), with the two alternative weightings.

			Corpus		Prominence under re-weighting		
Rank	**Country**	**BPI**	**Projects**	**Sectors**	**Region**	**BPI count-heavy**	**BPI breadth-heavy**
1	EGY	0.92	90	10	North Africa	0.93	0.92
2	PAK	0.83	69	7	South Asia	0.82	0.83
3	KEN	0.83	75	8	Sub-Saharan Africa	0.83	0.82
4	ETH	0.78	71	7	Sub-Saharan Africa	0.78	0.78
5	KHM	0.70	56	6	Southeast Asia	0.69	0.70
6	LAO	0.69	51	5	Southeast Asia	0.67	0.69
7	SAU	0.69	36	6	Middle East	0.63	0.70
8	SRB	0.68	53	4	Central-Eastern Europe and Balkans	0.66	0.67
9	RUS	0.68	49	5	CIS	0.65	0.67
10	TZA	0.68	46	7	Sub-Saharan Africa	0.64	0.68
11	NGA	0.67	20	7	Sub-Saharan Africa	0.58	0.69
12	ZAF	0.67	12	6	Sub-Saharan Africa	0.56	0.69
13	ARE	0.67	26	7	Middle East	0.59	0.69
14	UGA	0.65	23	7	Sub-Saharan Africa	0.57	0.67
15	BRA	0.65	24	6	Latin America and the Caribbean	0.57	0.66
16	DZA	0.64	26	5	North Africa	0.57	0.65
17	PER	0.64	14	5	Latin America and the Caribbean	0.54	0.65
18	ZMB	0.64	29	7	Sub-Saharan Africa	0.57	0.66
19	ARG	0.64	25	6	Latin America and the Caribbean	0.56	0.65
20	LKA	0.64	22	6	South Asia	0.56	0.65
Host countries			140				
Canonical projects			1,688				

Note. BPI is the equal-weighted mean of four min-max scaled components; the two right-hand columns re-weight them.

The sectoral mix shifts toward project types that travel well to foreign audiences (Fig. 4). Transport remains largest and hard-connectivity still dominates at 90.8 percent, but the routine civil construction that fills the internal projection almost vanishes outward (16.2 percent → 0.2 percent) while social and community projects rise roughly

forty-five-fold (to 18.0 percent). The outward projection is not a softer BRI; it is a selectively foregrounded, flagship-amplified version of the same hard-infrastructural backbone. By symbolic salience we mean projects with high public visibility relative to their operational footprint (flagship corridors, signature buildings, multilingually-amplified showcases), and this defines the strategic or showcase category used below: projects carried across multiple outward languages and/or projected at a markedly higher rank than their internal footprint warrants.

Table 5
Sector composition of the internal (Map A) and outward (Map B) projections.

Sector	**Internal projection (Map A)**		**Outward projection (Map B)**	
	Internal n	**Internal %**	**Outward n**	**Outward %**
transport	1246	32.0	656	38.9
energy	1166	30.0	355	21.0
civil construction	631	16.2	3	0.2
industrial	288	7.4	200	11.8
water	202	5.2	0	0.0
mining	71	1.8	0	0.0
education	67	1.7	0	0.0
agriculture	65	1.7	65	3.9
health	61	1.6	2	0.1
environment	37	1.0	0	0.0
multi sector	20	0.5	34	2.0
social	17	0.4	304	18.0
ict	16	0.4	30	1.8
other	5	0.1	1	0.1
trade	0	0.0	38	2.3
Total	3892	100.0	1688	100.0
Projects			3,892 / 1,688	

Note. Shares are of canonical projects in each layer and may not sum to 100 because of rounding. Hard-connectivity projects, reported as a share in the text, are those carrying the portal's explicit □□□ designation plus those whose sector defaults to hard, so the hard share exceeds the sum of the six infrastructure rows.

Multilingual amplification confirms this picture. Most outward projects appear in only one language (95.3 percent; mean 1.06 languages per project); genuine multi-language amplification is reserved for a handful of flagships: the China–Russia East-Route Gas Pipeline (all five languages), the China–Laos Railway (four), and the Jakarta–Bandung High-Speed Railway, China–Pakistan Economic Corridor, Pelješac Bridge and Colombo Port City (three each). The outward projection is thus a selective amplification structure rather than a translated mirror of the internal field, broadcasting flagship projects across audiences while most projects remain single-language and audience-specific.

Table 6
Language amplification across the five outward portals.

Languages per outward project	Projects	Share (%)
1	1608	95.3
2	66	3.9
3	12	0.7
4	1	0.1
5	1	0.1
Total	1688	100.0
Outward projects	1,688	

Note. Languages per outward project, after consolidation on shared portal project identifiers.

The regional redistribution is equally legible. Sub-Saharan Africa remains the largest share (31.8 percent across forty-five hosts), but Southeast Asia falls 9.6 points to 13.3 percent, North Africa rises from 3.8 to 8.5 percent, and the CIS rises from 1.0 to 4.9 percent on the strength of Russia and Belarus, while Oceania nearly vanishes (1.8 to 0.5). Outward concentration is also slightly higher than internal (the top ten countries take 35.3 percent of canonical projects against 32.2 internally), consistent with the stronger narrative selectivity expected of the outward register.

Table 7
Regional composition and mean prominence in the two projections.

Region	Internal projection (Map A)				Outward projection (Map B)				
	Internal projects	Internal %	Internal hosts	Internal mean BPI	Outward projects	Outward %	Outward hosts	Outward mean BPI	Δ share (pp)
Sub-Saharan Africa	1197	30.8	46	0.51	536	31.8	45	0.48	+ 1.0
Southeast Asia	891	22.9	11	0.69	224	13.3	11	0.56	-9.6
Middle East	387	9.9	11	0.52	156	9.2	12	0.54	-0.7
Latin America and the Caribbean	331	8.5	23	0.50	149	8.8	23	0.49	+ 0.3
South Asia	326	8.4	7	0.57	160	9.5	7	0.55	+ 1.1
Central Asia	221	5.7	5	0.64	78	4.6	5	0.54	-1.1
Central-Eastern Europe and Balkans	185	4.8	11	0.52	101	6.0	11	0.49	+ 1.2
North Africa	148	3.8	4	0.56	143	8.5	4	0.65	+ 4.7
Oceania	71	1.8	11	0.38	8	0.5	4	0.46	-1.4
East Asia	40	1.0	2	0.50	18	1.1	3	0.45	+ 0.0
CIS	39	1.0	4	0.48	82	4.9	3	0.56	+ 3.9
Western Europe	35	0.9	11	0.43	23	1.4	9	0.48	+ 0.5
Caucasus	19	0.5	3	0.53	9	0.5	2	0.44	+ 0.0
North America	2	0.1	2	0.50	1	0.1	1	0.38	+ 0.0
Total	3892	100.0	151	—	1688	100.0	140	—	—
Host countries			151 / 140						

4.4 Divergence

The divergence between the two projections is the article's central result, and three quantitative findings frame it. First, the association is moderate, not tight: across the 133 countries in both projections, internal and outward project counts correlate at r = 0.67, BPI at r = 0.55, and the rank correlation between BPI scores is ρ = 0.61, too

positive to be unrelated, too far from unity to be translation. When the two projections are set against a coverage-matched null (Section 4.6), the divergence is far larger than independent sampling of a single latent prominence, at the two corpus sizes, would produce, so the gap is genuine selectivity, not sampling noise.

Second, the rank shifts are large and patterned. The sharpest outward over-projections relative to internal standing are Djibouti (rank 141 → 23), Belarus (135 → 45), Czechia (150 → 64), the Netherlands (147 → 64), France (130 → 49), Germany (136 → 59), Kuwait (100 → 25), Qatar (95 → 30), Spain (98 → 39) and Tunisia (114 → 55); the sharpest dampings are the Philippines (15 → 102), El Salvador (56 → 133), Ghana (24 → 91), Romania (69 → 134), Benin (57 → 120), India (63 → 120), Tajikistan (40 → 96), Gabon (65 → 120), Mauritania (48 → 102), Bolivia (39 → 87), Georgia (58 → 104), Mongolia (41 → 85) and Vietnam (35 → 79). The over-projected cluster is dominated by Gulf showcase states and Western European rail termini whose outward narrative emphasizes strategic or showcase projects (Djibouti's corridor, Gulf energy and industrial parks, the China–Russia gas-pipeline coverage that also lifts Belarus, terminus references on the China–Europe Railway Express) against a near-zero domestic-reporting base, at intensities beyond their routine footprint. The damped cluster mixes strategically sensitive partners (the Philippines, India and Vietnam among them) with operationally substantial but lower-profile African and Latin American hosts (Gabon, Ghana, Bolivia, Mongolia) whose large internal footprints barely carry over outward.

Table 8
The ten sharpest outward over-projections and the ten sharpest dampings, by rank over all 158 union countries.

Country	Internal rank	Outward rank	Shift
Panel A. Ten sharpest outward over-projections			
DJI	141	23	+ 118
BLR	135	45	+ 90
ATG	152	64	+ 88
PRK	152	64	+ 88
SVN	152	64	+ 88
CZE	150	64	+ 86
NLD	147	64	+ 83
FRA	130	49	+ 81
DEU	136	59	+ 77
KWT	100	25	+ 75
Panel B. Ten sharpest outward dampings			
PHL	15	102	-87
SLV	56	133	-77
BGR	73	141	-68
GHA	24	91	-67
ROU	69	134	-65
BEN	57	120	-63
MDA	80	141	-61
BHS	82	141	-59
AUT	82	141	-59
ARM	82	141	-59
Union countries ranked	158		

Note. Ranks run over all 158 union countries; a country absent from one layer takes the lowest rank there.

The strategically sensitive members of that damped cluster point to a mechanism the design can surface but not isolate. For partners managing open frictions with Beijing, conspicuous BRI branding carries a reputational cost on the host side, not only on China's: the Philippines and Vietnam calibrate the visibility of their alignment amid maritime and South China Sea disputes (Kuik 2016); India, a non-participant that declined to join and boycotted the 2017 Belt and Road Forum over a corridor running through territory it claims, is a limiting case whose low outward prominence is over-determined and carries no selection signal. Muted outward prominence for the hedging partners may therefore register their own calibration as much as China's editorial selection. The damped cluster is, in this

sense, not uniform: high-footprint hosts such as Gabon and Mongolia fall outward for want of narrative amplification, whereas the sensitive partners fall for reasons that plausibly originate with the host, a distinction the dual-projection design exposes but, lacking host-side data, cannot adjudicate.

Third, and most tellingly, Indonesia and Egypt are near-inverses. Indonesia leads Map A (0.92) but falls to 0.60 outward; Egypt is mid-tier internally (0.57) yet tops Map B (0.92), strong outward amplification atop a moderate operational footprint. In the Indonesian case the divergence is cleanest, and it survives the most demanding control. Scaling prominence against the size of the host economy, internal and outward prominence each track GDP only moderately (Spearman $\rho \approx 0.58$): bigger economies are on average more prominent, but the residuals are the finding. Indonesia, among the largest economies of any confirmed host (fifth by GDP), sits first in the internal projection, slightly above even its economic weight, yet 26th outward, under-projected by twenty-one ranks relative to its GDP. The residuals partition cleanly: strategically sensitive or large economies (Iran, India, the Philippines, Vietnam, EU members) are under-projected outward, while small economies carrying flagship corridors, Laos, Cambodia, Djibouti and Rwanda among them, are over-projected. The outward damping cannot, then, be explained away as proportionality to economic size. Indonesia's outward profile rests heavily on a single flagship, the Jakarta–Bandung high-speed railway, carried across several languages, while the dense long tail of its domestic-facing portfolio does not travel outward at proportional intensity. Egypt is where outward prominence most clearly does work the operational record cannot. The amplification settles on a node whose value to a connectivity narrative lies in its geography. Egypt holds the Suez–Mediterranean–Red Sea pivot and reads simultaneously as Arab, African and Mediterranean. Foregrounding it lets the outward projection present BRI as a coherent trans-regional space across those seams, a coherence the internal map, concentrated in Southeast Asia and Sub-Saharan Africa, does not itself supply. This is the clearest case of the region-making described in Section 2.2: the outward cartography narrates a region around a hub rather than mirroring where projects actually cluster. That the amplification holds across all five portals, not only the Egypt-centered Arabic stream, shows it is not an artifact of a single audience.

Table 9
Ranks and residuals net of economic size, selected hosts.

Country		Internal projection		Outward projection	
	GDP rank	Internal rank	Internal residual	Outward rank	Outward residual
IDN	5	1	+ 4	26	-21
KHM	55	5	+ 50	5	+ 50
LAO	84	7	+ 77	6	+ 78
DJI	103	111	-8	21	+ 82
RWA	89	55	+ 34	22	+ 67
IRN	11	115	-104	114	-103
IND	1	57	-56	97	-96
PHL	13	14	-1	81	-68
VNM	14	32	-18	61	-47
ROU	19	64	-45	104	-85
Hosts matched to GDP			121		

Note. Residuals are GDP rank minus prominence rank; negative outward residuals indicate under-projection net of size.

If read together, these results support three interpretations the design makes available without assuming them. The two projections are correlated but distinct: outward publication is selective amplification, not translation. The selectivity is structured, favouring some regions (North Africa, the CIS, the Balkans) and countries (Egypt, the Gulf, Djibouti) while damping others (Indonesia, Malaysia, the Philippines, Vietnam); its clearest single signature is the Southeast-Asia-for-North-Africa-plus-CIS trade-off (− 9.6 against + 8.5 percentage points). And the hard-connectivity backbone is preserved across projections even as the soft-connectivity composition shifts. Whether this curated prominence also corresponds to material engagement is the question Section 4.7 takes up.

4.5 Does audience specialization explain it?

Map B aggregates five portals whose audiences differ, so the apparent selectivity might be mechanical. Each non-English portal does over-index its home region. The Spanish portal places 53 percent of its projects in Latin America, against 9 percent in Map A, led by Brazil, Argentina, Peru, Chile and Ecuador. The Russian portal places 38 percent in the CIS and Central Asia, against 8 percent, led by Russia, Kazakhstan and Uzbekistan. The Arabic portal foregrounds the Middle East and North Africa (Egypt, the UAE, Saudi Arabia), the French portal francophone Africa, and the English portal serves the most diffuse readership and stays closest to Map A.

Table 10
Cross-layer Spearman ρ versus internal Map A, by outward language portal.

Portal			Cross-layer association			
	Canonical projects	Host countries	n matched	ρ vs Map A	95% interval	Top country
English	749	103	100	0.658	0.515 to 0.767	Pakistan
Arabic	319	73	71	0.054	−0.186 to 0.283	Egypt
French	265	73	72	0.468	0.254 to 0.642	Kenya
Russian	235	65	64	0.523	0.306 to 0.689	Russia
Spanish	217	66	64	0.316	0.089 to 0.509	Brazil
Portals			5			

Note. n is the number of host countries the portal carries; rho is against the internal BPI on those countries.

Taking that audience baseline as given, the diagnostic question is what deviates from it, and two patterns deviate across every portal, not merely within a home audience. Indonesia is damped in all five (ranking 30/11/39/10/19 in the English/Arabic/French/Russian/Spanish portals, top-ten only in the Russian portal despite leading Map A), and Egypt is amplified in all five (12/1/3/12/7), most sharply in Arabic but not only there; the Southeast-Asian share is negative against Map A in every portal, from − 4 points in English to − 15 in French. Audience specialization explains the regional baselines; it does not explain these cross-language consistencies. Were Beijing merely adapting to each audience's natural regional interests, the same countries would not be damped or amplified across all five streams. They are, and that is the signature of genuine outward selectivity. The sets around the two headline cases behave the same way. Malaysia, in the internal projection's second tier (0.80), ranks 26th, 19th, 45th, 20th and 14th across the English, Arabic, French, Russian and Spanish portals, and the Philippines and Vietnam drop out of the French and Russian portals altogether; Pakistan, at 0.74 internally against Indonesia's 0.92, ranks first, fourth, seventh, third and sixth. Kenya and Russia, by contrast, are lifted only where a home audience is served, Kenya in the English and French portals and Russia in the Russian, which is audience specialization at work, the baseline the cross-language cases depart from.

While the English portal stays closest to the internal map (ρ = 0.66), the Arabic portal is essentially uncorrelated with it (ρ = 0.05), constructing an almost independent cartography centered on Egypt, Pakistan, Saudi Arabia and Ethiopia. The decomposition thus isolates two phenomena: ordinary audience specialization, in which each portal foregrounds its own region, and a residual cross-language selectivity that audience tailoring cannot explain, and it is the second that the consistent Indonesia damping and Egypt amplification expose.

4.6 Robustness

Three checks confirm the divergence is not an artifact of construction. Under count-heavy and breadth-heavy alternatives to the headline BPI, country rankings correlate with the headline ranking at ρ ≥ 0.989 in both projections, and the Indonesia and Egypt divergences retain their sign and magnitude. The coverage-matched null preserves country structure: each country's pooled internal-and-outward project total defines a single latent prominence, from which an internal sample (3,441 projects) and an outward sample (1,498 projects), the within-125 subtotals of the full 3,892 internal and 1,688 outward canonical projects, the 125 mutually-confirmed countries being those on which per-country counts allow the simulation, are drawn independently as multinomial samples and the cross-projection

rank correlation recomputed, 3,000 times. If the two projections were only different-sized samples of one underlying prominence, the null places their correlation at 0.93 (95 percent interval 0.91 to 0.95). The observed correlation across those same 125 countries falls below the entire null distribution ($p < 0.001$), and it does so on the statistic the null simulates: the rank correlation of the two project counts is 0.79 against a null band of 0.91 to 0.95. The BPI rank correlation over the same countries, 0.69, lies further below again; the headline 0.61 is the full 133-country dual-projection set. What the null rules out is a specific alternative: that the two projections are independent samples of one latent prominence at these corpus sizes. It does not rule out every mechanical account. The two corpora differ in producers, in genre and in time window, and any such difference would also lower the observed correlation below the null band. The multinomial draw is itself an idealisation, since flagship amplification makes real outward sampling lumpier than the multinomial assumes, and a lumpier process would lower the null correlation and narrow the gap. The divergence is therefore established against sampling at equal prominence, and the question of which production difference produces it is left open in Section 5. The weighting checks are equally reassuring in detail. Under a count-heavy scheme (0.40 on project count, 0.20 on each remaining component) and a breadth-heavy scheme (0.20/0.30/0.25/0.25), the top-twenty country set overlaps the headline ranking by eighteen to twenty of twenty in both projections, and the headline divergences keep their sign and scale. Indonesia's internal-minus-outward gap is + 0.33, + 0.41 and + 0.31 across the three weightings, Egypt's − 0.35, − 0.41 and − 0.34. No weighting choice manufactures or erases the divergence. Crediting cross-border projects to every named host is a construction choice, so the consolidation was rerun under a stricter rule that assigns each of them only to its primary host, the first non-China country in the source record. The regional shift survives: Southeast Asia still falls 9.7 points, North Africa still rises 5.0 and the CIS 4.3, and the cross-layer correlation strengthens slightly, from 0.61 to 0.65. The rule's main effect is to remove five Western states from Map A altogether (Austria, Switzerland, Iran, the Netherlands and Czechia), which had appeared only where cross-border lines terminate.

Table 11
Robustness of the divergence across weighting, null and allocation rules.

Check	Statistic	Result
Headline cross-projection ρ (133 countries)	Spearman ρ	0.61 [0.49 to 0.71]
Count-heavy weighting	ρ vs headline, internal / outward	0.996 / 0.989
Breadth-heavy weighting	ρ vs headline, internal / outward	0.998 / 0.995
Top-twenty overlap under re-weighting, internal	count-heavy / breadth-heavy	19 / 18 of 20
Top-twenty overlap under re-weighting, outward	count-heavy / breadth-heavy	18 / 20 of 20
Indonesia gap	headline / count-heavy / breadth-heavy	+ 0.33 / +0.41 / +0.31
Egypt gap	headline / count-heavy / breadth-heavy	-0.35 / -0.41 / -0.34
Coverage-matched null	null ρ (95% interval)	0.93 (0.91–0.95)
Observed ρ, same 125 countries	Spearman ρ	0.69 ($p < 0.001$)
Primary-host-only consolidation	cross-layer ρ	0.65
Dual-projection countries	133	

Note. The null draws N = 3,441 internal and N = 1,498 outward projects from one latent prominence.

These reweighting results follow from the index's internal structure. The ranking tracks mainly project count and sector diversity, themselves correlated at 0.93, with which the composite correlates at 0.79 and 0.72 respectively; hard-connectivity share contributes little (0.18, interval spanning zero) and path-dependence moderately (0.39). The index is therefore best read as a count-and-breadth measure with a modest path-dependence adjustment, and the reader should take that as the honest description of what it is: a composite whose ranking is dominated by project count, with sector diversity adding little independent information because it is largely a function of count. Its four components are not independent dimensions, and the reweighting invariance above ($\rho \geq 0.989$) is a consequence of this collinearity rather than independent corroboration. The ranking is stable because it is substantially one underlying quantity, not because three further components independently corroborate it. We retain the composite over raw project count because the breadth and path-dependence components carry interpretively relevant information about the kind of engagement a country hosts, even where they do not alter its rank.

Table 12
Correlation structure among the Prominence Index components.

Pair	ρ	95% interval
BPI × project count	0.785	0.691 to 0.852
BPI × sector diversity	0.725	0.622 to 0.805
BPI × hard-connectivity share	0.183	−0.001 to 0.357
BPI × path dependence	0.394	0.220 to 0.542
Project count × sector diversity	0.929	0.903 to 0.946
Internal hosts	151	

4.7 Does prominence track material engagement?

The maps so far show how Beijing curates BRI space. Whether that curation tracks material engagement is a separate question. If the portal corpus were communicative emphasis without underlying activity, we would expect weak overlap with an independent record of Chinese finance and weak rank correlation with it. We test this against AidData's BRI-Tagged dataset (Burgess and Custer 2022). The tag is AidData's own judgement, independent of Beijing's labelling: a project is coded BRI-Like when its documented activity matches the five thematic priorities of Xi's 2013 speeches, whether or not any Chinese source ever branded it as Belt and Road, and the affirmative set runs to 4,723 projects across 138 host countries between 2000 and 2017. The overlap is informative for that reason. A match between the portal and a list of projects Beijing itself had labelled BRI would be close to tautology; a match with a record built by other hands, from other documents, under a thematic rule, is evidence.

The overlap is substantial. Of the 158 portal-corpus countries, 125 (79.1 percent) are also AidData BRI-Like hosts, and 90.6 percent of AidData BRI-Like countries appear in the portal corpus, convergence between two independently-constructed methodologies, one from Beijing's curated representation, the other from external finance records. The combined universe is 171 countries, of which 73 percent appear in both records. More demanding is the rank correlation, and here economy size has to be partialled out before the association can be read as anything but scale: portal prominence and AidData finance both track GDP, so the test that matters is the association net of it.

In full, across the 121 of these hosts matched to World Bank GDP data, the raw rank correlation between portal prominence and AidData project count is 0.377 for the internal projection and 0.371 for the outward, 0.404 pooled; against adjusted USD finance, 0.600 and 0.593. Net of host GDP, these USD correlations fall to roughly 0.43, so a substantial share of the raw association reflects shared dependence on economy size; what survives the adjustment is moderate and positive, well short of a tight correspondence. Since the USD correlations run well above the count correlations, portal prominence aligns more closely with the financial scale of engagement than with the number of projects, which is what to expect if curation foregrounds high-value flagship engagement, the reading Section 4.4 gave the portal architecture on qualitative grounds.

Table 13
Association between portal prominence and AidData's finance record (n = 121)

Measure	Internal		Outward	
	Internal ρ	95% interval	Outward ρ	95% interval
AidData project count, raw	0.377	0.207 to 0.528	0.371	0.201 to 0.525
AidData USD finance, raw	0.600	0.460 to 0.715	0.593	0.440 to 0.720
AidData USD finance, net of host GDP	0.438	—	0.429	—
Hosts matched to GDP			121	
Note. n = 121 hosts matched to World Bank GDP; partial rho controls for log GDP.				

A ranked mutual-confirmation index (Map C, Fig. 5) re-orders the 125 confirmed hosts by a joint evidence score, the mean of z-standardised portal prominence and z-standardised log AidData project count. Its top is held by countries prominent on both records at once: Indonesia (joint score 1.70), Cambodia (1.66), Pakistan (1.63), Ethiopia (1.60), Laos (1.52), Angola (1.48), Kenya (1.47), Tanzania (1.31), Zambia (1.16) and the Democratic Republic of the Congo (1.07), with Bangladesh and Myanmar close behind. Indonesia's place at the head of this list has to be squared with Section 4.4, where the same country is the clearest case of outward damping, and the two results describe two axes. Indonesia is the most deeply embedded host in the corpus, first in Map A by a clear margin and among the top three AidData recipients by project count, and that embedding is what the joint score registers. Beijing's outward narration of the country is nevertheless muted in every language portal, and that is what Section 4.4 reports. On the audience axis Indonesia is under-projected; on the cross-method axis it is the best-confirmed host of all. A single-source map would have shown one or the other. The dual-projection design is what makes both visible at once.

Table 14
The ten highest-ranked mutually-confirmed hosts (Map C).

Rank	Country	Evidence score	Internal BPI (projects)	Outward BPI (projects)	AidData projects
1	IDN	1.70	0.92 (207)	0.60 (23)	119
2	KHM	1.66	0.79 (115)	0.70 (56)	119
3	PAK	1.63	0.74 (107)	0.83 (69)	91
4	ETH	1.60	0.73 (88)	0.78 (71)	102
5	LAO	1.52	0.78 (114)	0.69 (51)	98
6	AGO	1.48	0.68 (69)	0.59 (13)	156
7	KEN	1.47	0.77 (105)	0.83 (75)	65
8	TZA	1.31	0.79 (118)	0.68 (46)	69
9	ZMB	1.16	0.67 (35)	0.64 (29)	82
10	COD	1.07	0.64 (47)	0.54 (13)	98
Mutually confirmed hosts			125		
Note. Evidence score is the mean of z-standardised portal BPI and z-standardised log AidData count.					

The disagreements are themselves informative, and legible in detail. The 33 portal-only countries are led by the Gulf (Saudi Arabia, the United Arab Emirates, Qatar, Kuwait and Bahrain), some fifteen Western and Central-Eastern European states (Germany, France, Italy, Spain, Poland, Greece, Hungary, Croatia and others), Singapore and South Korea, and a few Pacific micro-states, engagements that are commercial, infrastructural or strategic, and so lie outside AidData's development-finance frame. What puts them on the portal is Gulf energy and industrial deals, European rail termini and bridge contracts, and commercial port stakes. The 13 AidData-only countries run the other way: Venezuela (36 BRI-like projects), Samoa (51), the Marshall Islands (25), Albania (18), Yemen (15), Syria (10), Haiti (10) and a handful more, several under sanctions or in protracted crisis, historical recipients the contemporary portal no longer foregrounds. What they share is muting under present political conditions: states in crisis or exposed to sanctions (Venezuela, Yemen, Syria, Libya, Somalia), Pacific recipients outside the active publicity programme, and Albania, whose engagement predates the current narrative architecture. The portal does not inherit the finance footprint mechanically; it chooses which past and present engagements to carry under the contemporary brand. Each disagreement marks a real difference between Chinese engagement abroad and its state-curated narration.

One scope condition applies to the comparison. AidData covers 2000–2017, whereas the portal corpus accumulates material spanning roughly 2015–2026 and concentrated in 2018–2025; the two windows therefore overlap only at their edges (2015–2017), with the portal corpus predominantly more recent. Some portal-only countries, in the Gulf, in Central and Eastern Europe and among recent African and Middle Eastern partners, may therefore be post-2017 hosts whose engagement AidData has had no occasion to record, in which case the category mismatch described above and a timing effect contribute jointly to the portal-only set. The near-disjoint windows mean the correspondence must be read cautiously: it partly reflects slow-moving country-level fundamentals (economy size and long-run engagement) common to both records, and to that extent does not validate contemporaneous curation, which is why the size-adjusted correlation above is the more conservative and appropriate test.

5. Conclusion

The article opened with an elementary question, where does the BRI begin and end, and answered it by refusing two unsatisfactory positions: treating the initiative as incoherent for lack of a fixed boundary, or accepting an inherited country list as if the boundary problem were solved. In place of membership counts, it puts prominence in two projections and a test against independently recorded finance. The central finding is that Beijing's two projections diverge in structured ways ($\rho = 0.61$). Indonesia and Malaysia damped in outward materials, Egypt and the Gulf amplified, Southeast Asian regional weight traded for North African and CIS visibility, a divergence consistent across all five outward languages and robust to scaling for economic size. Two supporting results stand on their own. The first is an operational scope for officially represented BRI space: the Portal supplies a 158-country record, ordered by prominence, in place of the inherited baskets the field has reused. The second is the cross-method validation: portal prominence is moderately associated, net of economy size, with AidData's independent record of Chinese official finance, indicating the maps are partly grounded in material engagement, not communicative selection alone.

These findings speak back to several literatures. Against work that treats the BRI as a bounded region, they show that even Beijing's own representation does not produce one; against work that finds it too amorphous to measure, they show that operationalisation is feasible when the source universe is held constant and the audience varied. In

the terms of the regionness literature invoked in Section 2.2, the maps give the BRI's contested regionhood an operational footing: how it is represented becomes measurable by degree.

Three limits bound what the maps license. First, the portal-authority assumption admits a long tail: a few countries appear with a single canonical entry (the United States and Canada among them), retained at face value without any external threshold. The consequence is modest, since seventeen Map A countries have one project and twenty-eight two or fewer, together about one percent of internal records, so the headline findings, which rest on countries with project counts in the dozens or hundreds, are unaffected. A floor effect follows from the same curation logic. When a country has one canonical project, no explicit hard-or-soft flag and no path-dependence code, it takes the sectoral defaults on both components, and min-max scaling then converts those defaults into a BPI near 0.50 that the record does not earn. Sweden, the United States and Canada all sit at exactly 0.50 in Map A for this reason. The article therefore anchors its interpretation in countries above the long-tail threshold and publishes every component (project count, sector diversity, hard-connectivity share, path-dependence) so that readers can rebuild the ranking under an alternative threshold or weighting; the count-heavy and breadth-heavy columns exist for that purpose.

Second, the two projections do not cover identical time windows: the internal projection leans toward recent operational activity while the outward projection preserves a longer communicative memory, so individual divergence magnitudes should be read as combined operational-and-communicative signals. Although the direction is robust to this asymmetry, specific country magnitudes are more sensitive to it.

Third, the method measures representation, not its reception or effects. A high BPI registers Beijing's official projection, not a host state's reciprocal endorsement, project quality, or popular reception; the AidData correlation bridges representation to the financial scale of engagement but not to a full causal account of leverage, which would require project-level financial-exposure data and process-tracing that the dual-projection design does not attempt. Divergence, finally, is consistent with several non-exclusive mechanisms (audience management, reputational filtering, temporal asymmetry, the differential availability of language-specific materials), which the design surfaces but does not adjudicate between.

None of these destabilises the divergence finding; each bounds what the results support. The full canonical project database, the per-country BPI components, and the consolidation rules are available so that the maps can be rebuilt, contested, or extended on shared terms.

The dual-projection design extends in several directions. A longitudinal version would compare projections across years. Closer integration with project-level finance data would tighten the bridge from prominence to engagement. And the projections may differ in genre as well as content (routine reporting versus flagship narration), a difference this design makes newly examinable. The procedure is also portable: any regional project whose official spatial scope is curated can be mapped the same way, and, as the differentiated-integration literature shows, even treaty-based regions such as the Eurasian Economic Union, the Gulf Cooperation Council, and the African Continental Free Trade Area have functionally porous boundaries that their formal membership does not exhaust, so the design's reach is not confined to informal initiatives. In each case the move is the same: separate domestic-facing from outward-facing projection and validate against it. The BRI's official cartography is two maps rather than one, and their divergence is itself part of the politics of the initiative. Mapping that structure is a precondition for serious comparative analysis of Chinese spatial influence, and, by extension, of any regionalism that curates its own scope.

Declarations

Competing interests

The author has no competing interests to declare that are relevant to the content of this article.

Funding

No funding was received for conducting this study.

Data availability

The corpus is compiled from the official BRI Portal ecosystem (yidaiyilu.gov.cn and its English, Arabic, French, Russian and Spanish portals), a public source, together with AidData's BRI-Tagged Chinese Official Finance Dataset and World Bank GDP figures. The derived datasets (de-duplicated project database, per-country Prominence Index components, language-stratified series, GDP-scaled residuals and mutual-confirmation scores) are available from the corresponding author on reasonable request. The consolidation is auditable in its parts. The supplementary files carry the canonical project database for both layers, the per-country index components and their scaled values, the Map 0 list, the divergence dataset, the language-stratified outward aggregates, the cross-border primary-host sensitivity outputs, and a manual-review log recording the six sectors that defaulted to other and the one hundred and forty-two short project names retained for inspection. Because every canonical record carries the URL or URLs from which it was consolidated, a later researcher can verify source attribution record by record. Country-name disambiguation across five non-Chinese languages, sector aggregation across sub-categories and cross-border project allocation were all decided under uncertainty; the rules sit in Section 3.3 and the decisions are recoverable from the files. Replication is imperfect because some manual adjudication is irreducible, but the chain from URL-level observation to a map-level statement can be walked end to end.

Declaration of generative AI use

In preparing this manuscript the author used a generative AI model, Claude Fable 5.1 (Anthropic), to scan the full text and flag passages that might be improved or revised, and to assist with paraphrasing, spelling, grammar, repetition, and reference checking. The model produced no claim, evidence, or interpretation of its own; every analytical and editorial decision was made by the author, and every revision was reviewed and carried out by the author.

Ethics approval

The study analyses publicly available official publications and secondary datasets. No human participants were involved and no ethical approval was required.

References

1. Ang YY (2023) Ambiguity and Clarity in China's Adaptive Policy Communication. China Q 257:20–37
2. Bharti MS, Kumari S (2024) China's Belt and Road Initiative in Southeast Asia and Its Implications for ASEAN-China Strategic Partnership. *Asian Review of Political Economy*3 (1): Article 22

3. Brady A-M (2008) Marketing Dictatorship: Propaganda and Thought Work in Contemporary China. Rowman & Littlefield, Lanham, MD
4. Brady A-M (2015) Authoritarianism Goes Global (II): China's Foreign Propaganda Machine. J Democracy 26(4):51–59
5. Brautigam D (2020) A Critical Look at Chinese 'Debt-Trap Diplomacy': The Rise of a Meme. Area Dev Policy 5(1):1–14
6. Burgess B, Custer S (2022) BRI-Tagged Chinese Official Finance Dataset, Version 1.0. AidData at William & Mary, Williamsburg, VA
7. Cai P (2017) *Understanding China's Belt and Road Initiative*. Lowy Institute Analysis. Lowy Institute for International Policy, Sydney
8. Cheng J, Zeng J (2023) Digital Silk Road' as a Slogan Instead of a Grand Strategy. J Contemp China 33(149):823–838
9. Chiyemura F, Gambino E, Zajontz T (2023) Infrastructure and the Politics of African State Agency: Shaping the Belt and Road Initiative in East Africa. Chin Political Sci Rev 8(1):105–131. 10.1007/s41111-022-00214-8
10. Christen P (2012) Data Matching: Concepts and Techniques for Record Linkage, Entity Resolution, and Duplicate Detection. Springer-, Berlin
11. de Soyres F, Mulabdic A, Murray S, Rocha N, Ruta M (2018) How Much Will the Belt and Road Initiative Reduce Trade Costs? *World Bank Policy Research Working Paper No*. 8614. Washington, DC: World Bank
12. Dollar D (2019) Understanding China's Belt and Road Infrastructure Projects in Africa. Global China series, Washington, DC: Brookings Institution
13. Dukalskis A (2021) Making the World Safe for Dictatorship. Oxford University Press, New York
14. Favell A, Guiraudon V (2009) The Sociology of the European Union: An Agenda. Eur Union Politics 10(4):550–576
15. Fligstein N (2008) Euroclash: The EU, European Identity, and the Future of Europe. Oxford University Press, Oxford
16. Good PI (2005) Permutation, Parametric, and Bootstrap Tests of Hypotheses, 3rd edn. Springer, New York
17. Hameiri S, Jones L (2015) Global Governance as State Transformation. Polit Stud 64(4):793–810
18. Haralambides H, Merk O (2020) *The Belt and Road Initiative: Impacts on Global Maritime Trade Flows*. International Transport Forum Discussion Paper No. 2020/02. Paris: OECD Publishing
19. He A (2019) The Belt and Road Initiative: Motivations, Financing, Expansion and Challenges of Xi's Ever-Expanding Strategy. CIGI Paper No. 225. Centre for International Governance Innovation, Waterloo, ON
20. Heilmann S, Perry EJ (eds) (2011) Mao's Invisible Hand: The Political Foundations of Adaptive Governance in China. Harvard University Asia Center, Cambridge, MA
21. Hettne B, Söderbaum F (2000) Theorising the Rise of Regionness. New Polit Econ 5(3):457–472
22. Hillman JE (2020) The Emperor's New Road: China and the Project of the Century. Yale University Press, New Haven, CT
23. Hurley J, Morris S, Portelance G (2018) *Examining the Debt Implications of the Belt and Road Initiative from a Policy Perspective. CGD Policy Paper 121*. Washington, DC: Center for Global Development
24. Jones L, Hameiri S (2021) Fractured China: How State Transformation Is Shaping China's Rise. Cambridge University Press, Cambridge

25. Katzenstein PJ (2005) A World of Regions: Asia and Europe in the American Imperium. Cornell University Press, Ithaca, NY
26. Kuik C-C (2016) How Do Weaker States Hedge? Unpacking ASEAN States' Alignment Behavior towards China. J Contemp China 25(100):500–514
27. Leutert W (2020) State-Owned Enterprises in Contemporary China. In: Bernier L, Florio M, Bance P (eds) The Routledge Handbook of State-Owned Enterprises. Routledge, London, pp 201–212
28. Li H, Wang Y (2025) Adaptive Governance at Work: How Chinese State Actors Respond to Overseas Pushback on the Belt and Road Initiative. China Q 262:445–461
29. Liu W, Zhang Y, Xiong W (2020) Financing the Belt and Road Initiative. Eurasian Geogr Econ 61(2):137–145
30. Malik A et al (2021) Banking on the Belt and Road: Insights from a New Global Dataset of 13,427 Chinese Development Projects. AidData at William & Mary, Williamsburg, VA
31. Maliszewska M, van der Mensbrugghe D (2019) *The Belt and Road Initiative: Economic, Poverty and Environmental Impacts*. World Bank Policy Research Working Paper No. 8814. Washington, DC: World Bank
32. Mansfield ED, Milner HV (1999) The New Wave of Regionalism. Int Org 53(3):589–627
33. Nardo M, Saisana M, Saltelli A, Tarantola S, Hoffmann A, Giovannini E (2008) Handbook on Constructing Composite Indicators: Methodology and User Guide. OECD Publishing, Paris
34. Nedopil Wang C (2023) Ten Years of China's Belt and Road Initiative (BRI): Evolution and the Road Ahead. Griffith Asia Institute and Green Finance and Development Center, FISF, Fudan University, Brisbane and Shanghai
35. Repnikova M (2022) Chinese Soft Power. Cambridge University Press, Cambridge
36. Rolland N (2022) China's Belt and Road Initiative – Strategy: China's Evolving Ambitions. In *China's Belt and Road Initiative: A Geopolitical and Geo-economic Assessment*, 14–39. IISS Strategic Dossier. London: International Institute for Strategic Studies
37. Russel DR, Berger B (2019) Navigating the Belt and Road Initiative. Asia Society Policy Institute, New York
38. Schimmelfennig F, Winzen T (2020) Ever Looser Union? Differentiated European Integration. Oxford University Press, Oxford
39. Söderbaum F (2016) Rethinking Regionalism. Palgrave Macmillan, Houndmills
40. Sims JP, Lee Y-T, Lee BTF (2025) New Chinese Economic Policy to Latin America? A QCA Approach to the Belt and Road Initiative. Chin Political Sci Rev 10(1):1–27. 10.1007/s41111-023-00244-w
41. Vangeli A (2020a) Belt and Road and China's Attempt at Region Building in Central-East and Southeast Europe. J Curr Chin Affairs 49(1):14–32
42. Vangeli A (2020b) China's Belt and Road in the Balkans in the Post-COVID-19 Era. *In IEMed Mediterranean Yearbook*2020, 201–204. Barcelona: European Institute of the Mediterranean
43. Wang H (2021) The Belt and Road Initiative Agreements: Characteristics, Rationale, and Challenges. World Trade Rev 20(3):282–305
44. World Bank (2019) Belt and Road Economics: Opportunities and Risks of Transport Corridors. World Bank, Washington, DC
45. Yan H, Sautman B (2023) China, Ethiopia and the Significance of the Belt and Road Initiative. China Q 257:222–247
46. Yang H, Van Gorp B (2023) A Frame Analysis of Political-Media Discourse on the Belt and Road Initiative: Evidence from China, Australia, India, Japan, the United Kingdom, and the United States. Camb Rev Int Affairs

36(5):625–651
47. Ye M (2020) The Belt Road and Beyond: State-Mobilized Globalization in China, 1998–2018. Cambridge University Press, Cambridge
48. Yilmaz S, Li B (2020) The BRI-Led Globalization and Its Implications for East Asian Regionalization. Chin Political Sci Rev 5(3):395–416. 10.1007/s41111-020-00145-2
49. Zehra SN et al (2025) Redefining Diplomatic Discourse on the Belt and Road Initiative: A Corpus-Based Sentiment Analysis. Asian J Comp Politics 11(3). 10.1177/20578911251366121
50. Zeng J (2017) Does Europe Matter? The Role of Europe in Chinese Narratives of 'One Belt One Road' and 'New Type of Great Power Relations'. JCMS: J Common Market Stud 55(5):1162–1176
51. Zeng J (2019) Narrating China's Belt and Road Initiative. Global Policy 10(2):207–216

Figures

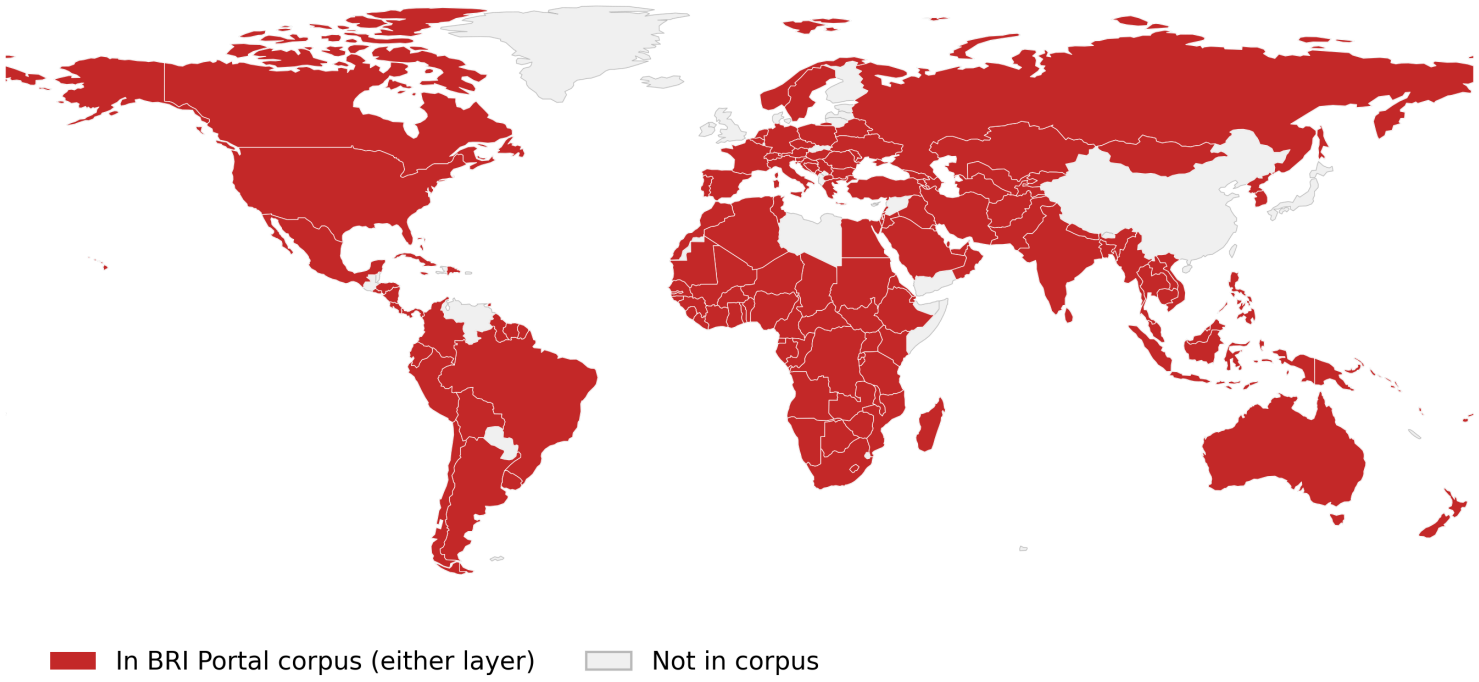


Figure 1

Map 0 Binary Portal Presence

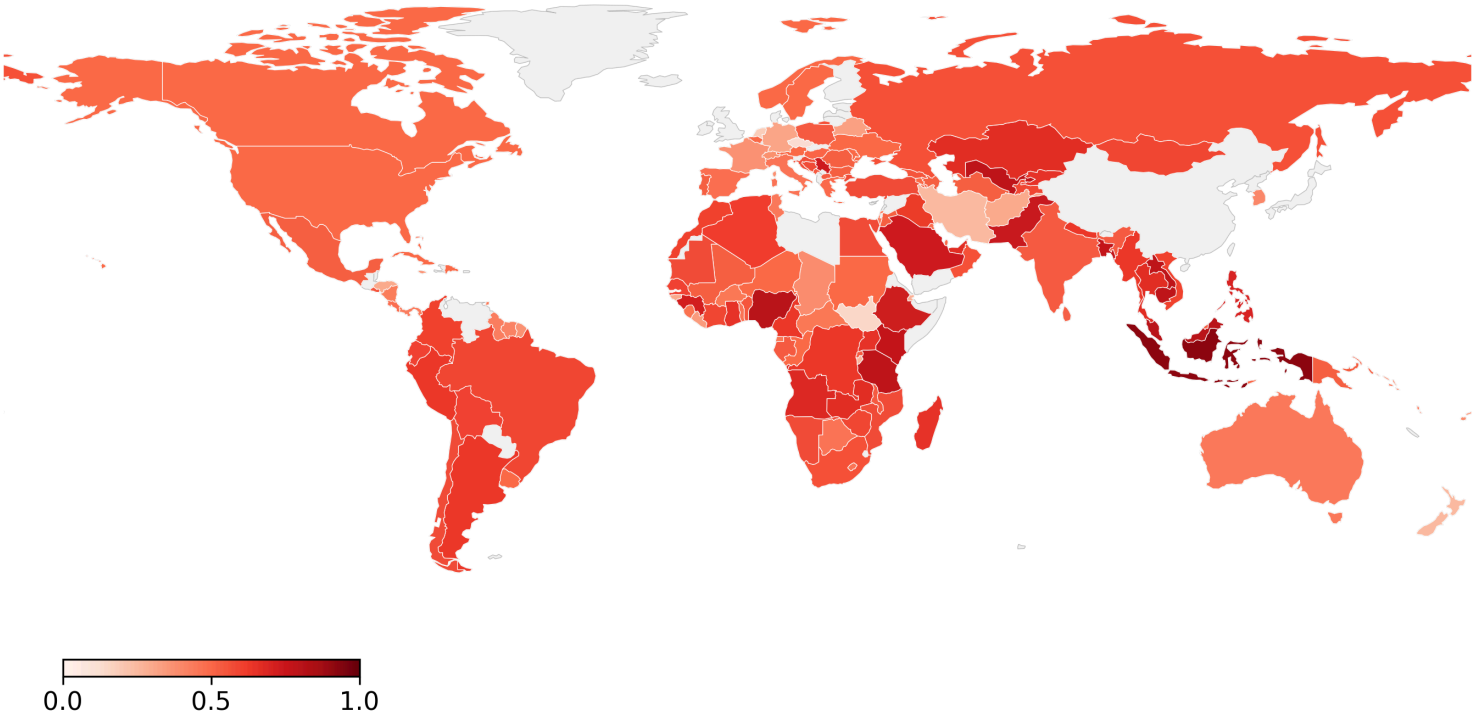


Figure 2

Map A Internal Projection

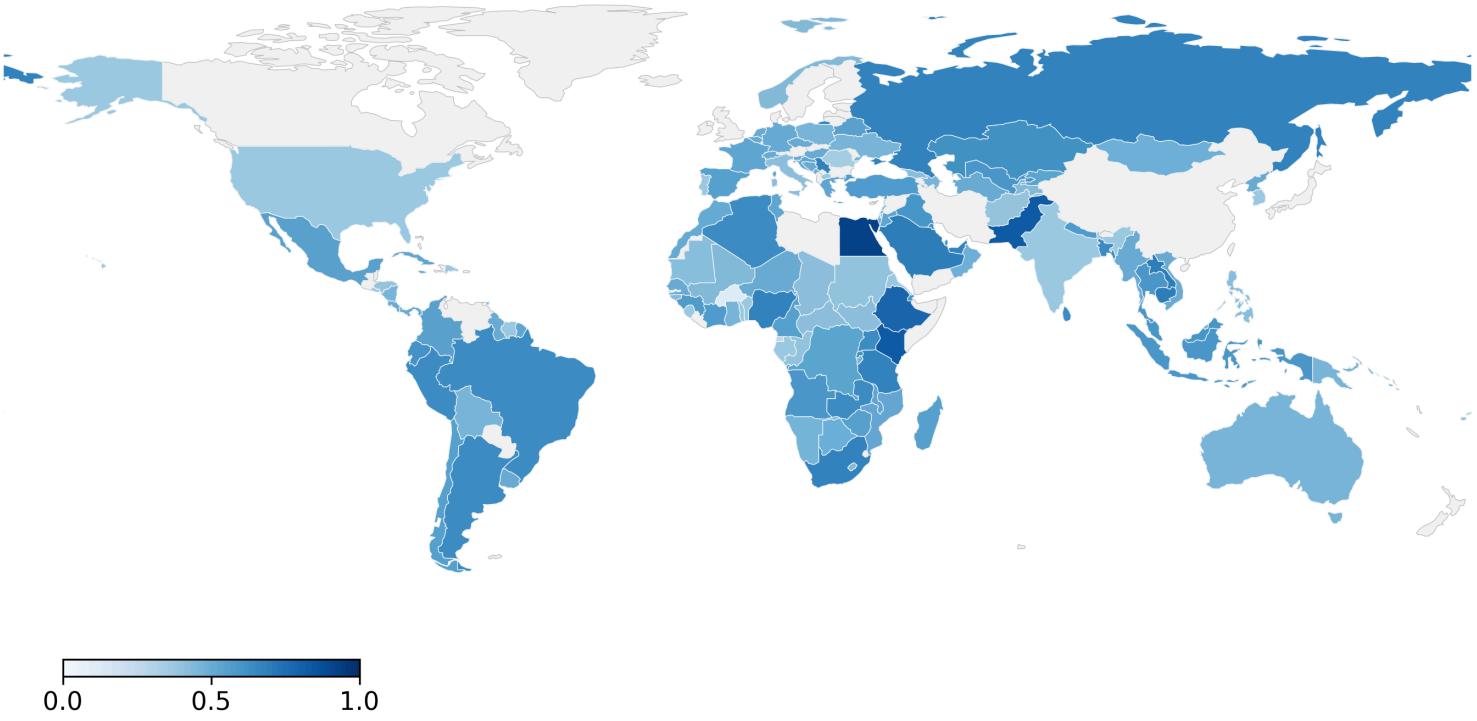


Figure 3

Map B Outward Projection

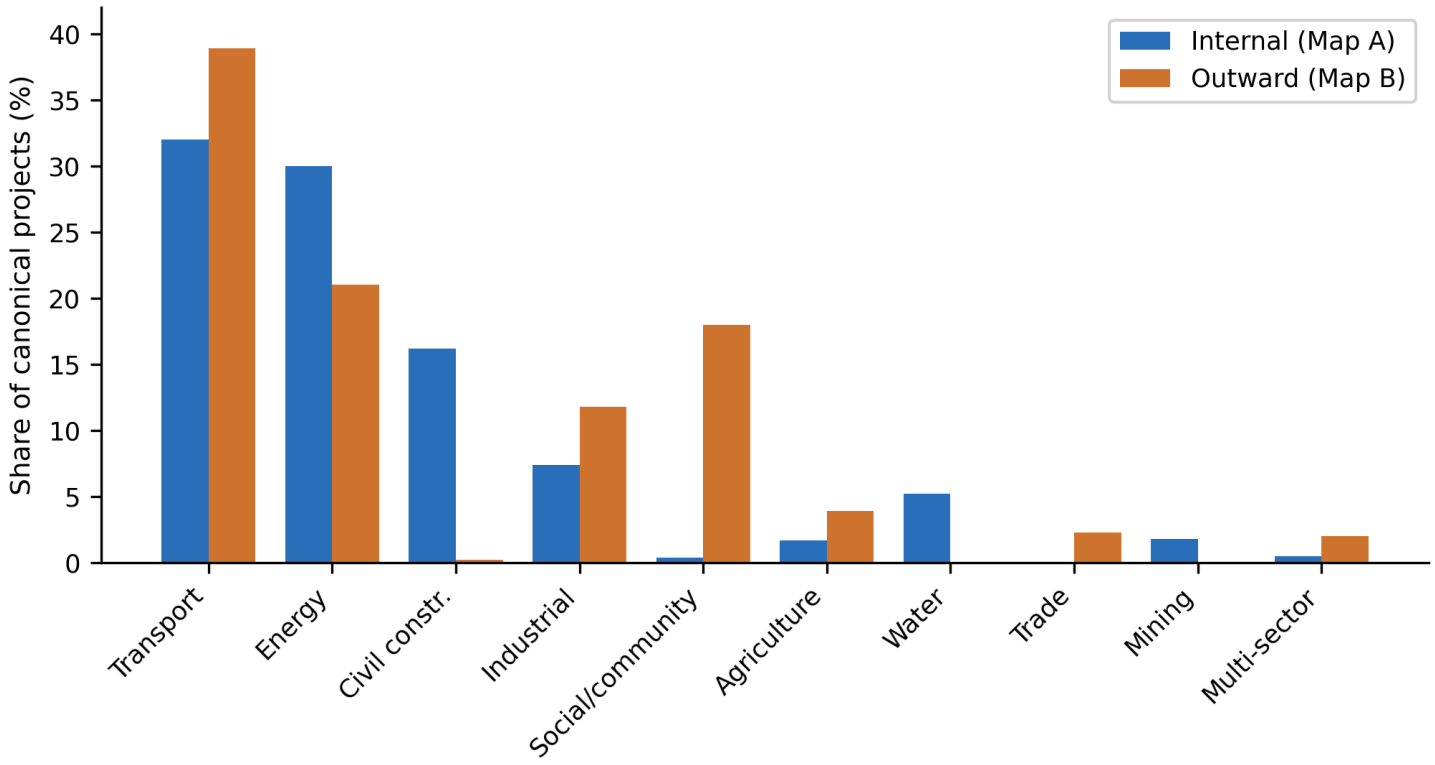


Figure 4

Sector Composition

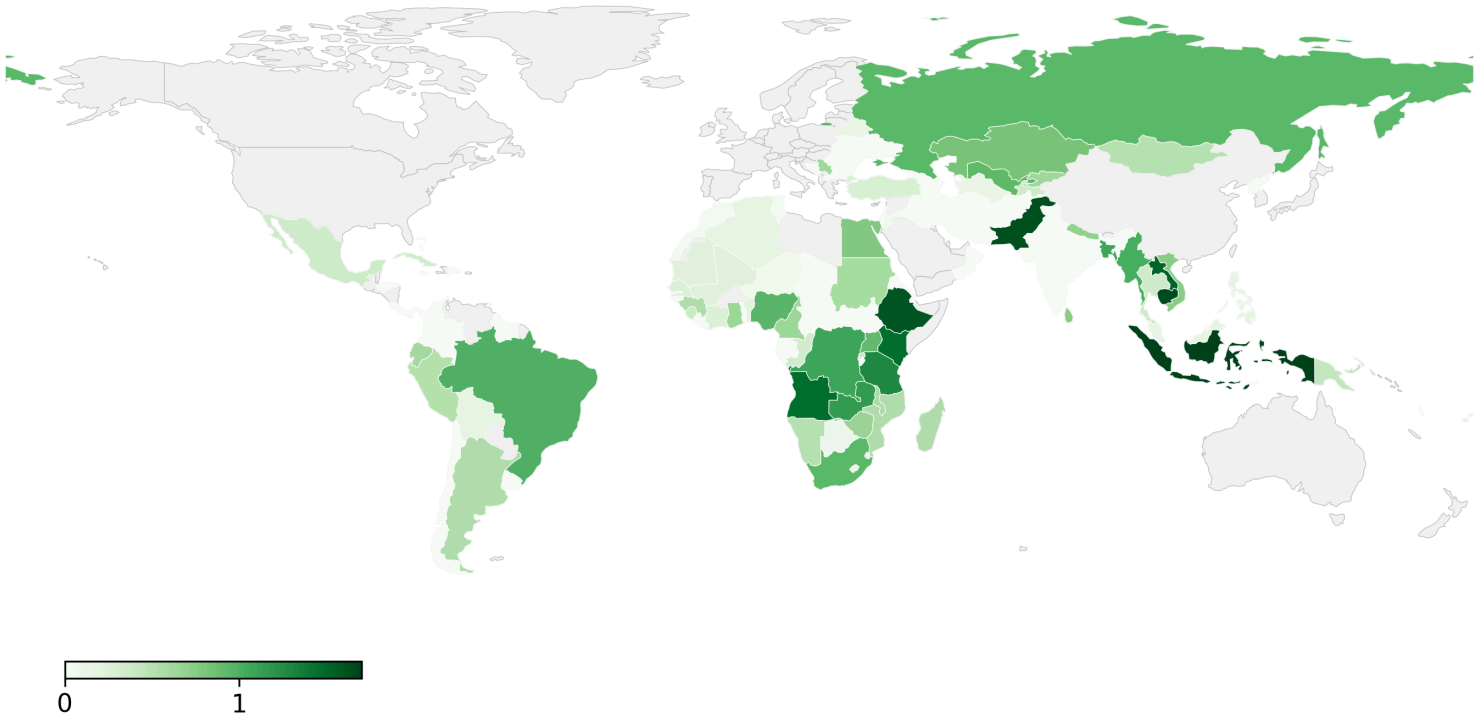


Figure 5

Map C Mutual Confirmation