

Capital Was Never the Constraint: The Missing Layer in Saudi Arabia's Innovation Economy

Thesis brief (v4) for a pre-submission opinion · adaptation for *Entrepreneurship & Regional Development*

Wael M. Altorabi

How to read this brief. It is written to be attacked, not to persuade. Every statistic carries a page reference (or a dataset identifier) and an evidentiary tier: [PR] peer-reviewed, [OI] official-institutional, [IP] industry-practitioner. Official figures read only through a press report are tagged [OI] but treated as lower-tier. † marks a figure from a primary document that is no longer in the project folder; it was recorded in the author's earlier coding matrix and must be re-verified. ‡ marks a page number taken from the project's page-referenced extraction of the RDIA and World Economic Forum report, whose original is not in the folder. Claims that cannot be independently verified are marked UNSUPPORTED rather than softened. The brief makes no recommendation concerning any firm or product.

1. Thesis and abstract

Saudi Arabia's state-led entrepreneurial ecosystem has accumulated venture capital and support programmes faster than it can demonstrate their conversion into firm outcomes, and its documentary record is best explained by a fragmentation equilibrium: each public agency is rewarded for activity it can attribute to itself, none can verify the value of integration, and a shared outcome layer is therefore never supplied.

Abstract. Entrepreneurial ecosystem theory locates performance in the relations among elements but treats the state as one framework condition. This brief examines Saudi Arabia, where public capital and programmes expanded rapidly after 2016 while measured outputs trail inputs. A qualitative analysis of eleven official documents, a research-authority programme report, a mapping of 418 support programmes and nineteen government data files identifies four patterns: abundance is documented but conversion is not; integrating mandates are stacked across agencies; measurement is vertical, with discordant and unattributable figures; and coordination is bilateral brokerage. These are read as a fragmentation equilibrium, stable because no agency gains from unilateral integration and integration is unverifiable. A UAE counter-case qualifies the argument, and refutation conditions are stated.

2. Research question and contribution

Research question. How do the public and quasi-public institutions of Saudi Arabia's entrepreneurial ecosystem describe, measure and coordinate support, and what do these documentary practices reveal about the governance of relations among ecosystem elements? The question is descriptive-explanatory. It does not ask whether fragmentation is the binding constraint on outcomes, which a document analysis cannot establish.

Position against the four reference frameworks. Isenberg's (2010) six domains (policy, finance, culture, support, human capital, markets) treat policy as one domain among six, so the relations *among* public bodies have no place in the model. Stam's (2015) systemic model, as operationalised by Stam and van de Ven (2021), makes elements interdependent and argues for a multiplicative rather than additive index (p. 826) [PR], but still assigns government "an enabling role ... in the background" (p. 813) [PR]. Their own regional data record that quality of government is "consistently negatively correlated with the other elements" (p. 820) [PR], an anomaly the framework notes but does not theorise. Spigel (2017) moves performance into the relations through which cultural, social and material attributes

reproduce one another, but assumes those relations are carried by networks and intermediaries in market-led settings; it does not ask who governs relations when the principal connectors are agencies with separate mandates. Sussan and Acs (2017) are the nearest neighbour: the digital entrepreneurial ecosystem framework brings the governance of digital infrastructure inside the ecosystem, and this brief's "shared layer" is a form of such infrastructure. The intended contribution relative to them is a political-economy account of why that infrastructure is under-supplied by the very public bodies that would have to build it, not an account of what it enables once present.

Contribution claimed. First, a governance level within relational ecosystem theory: in state-led ecosystems the relations among public agencies are part of the ecosystem's relational structure. Second, a mechanism, the fragmentation equilibrium, in which the absence of shared outcome measures is both a symptom and the stabiliser of fragmentation. Third, a method for reading official documents as evidence about the institutions that produce them.

Novelty is not yet cleared. The full text of Sussan and Acs (2017) is not held in the project corpus, so the characterisation above must be checked. A systematic review of 104 works argues that the ecosystem concepts are "old wine in new bottles" relative to the territorial approach of clusters and regional innovation systems (Scaringella & Radziwon, 2018, p. 1) [PR], so the governance level must be shown to add something beyond that older literature. The stronger threat comes from outside the four frameworks: public-administration research on whole-of-government coordination (Bouckaert et al., 2010; Christensen & Lægreid, 2007), digital public infrastructure and institutional voids. Question 2 in Section 8 turns on this.

3. Core argument: fragmentation as a stable equilibrium

The four documentary patterns could be read as administrative shortcomings awaiting better management. The brief argues instead that they describe a configuration that is stable because it is individually rational for each participant. The logic runs in five steps.

Step 1 (domain mandates). Public bodies were created between 2016 and 2021 with mandates defined by domain (SMEs, venture finance, investment, research, a sector) rather than by the founder's journey, so a typical founder crosses several mandates, and each new body acquires an integrating mandate for its own field. **Step 2 (attributable accountability).** Each agency is accountable for metrics it can attribute to itself: beneficiaries, sessions, capital committed. Cohort outcomes such as survival or follow-on financing are joint products of several programmes and cannot be attributed to any one. **Step 3 (asymmetric cost).** Contributing records to a shared layer, adopting common identifiers or referring clients onward imposes private costs on the contributing agency while the benefits accrue to founders and to the system, the classic structure in which a common good is under-supplied (Olson, 1965). **Step 4 (unverifiability).** Without a shared outcome layer, no agency can demonstrate that integrating improved outcomes, so integration goes unrewarded and coordination is supplied instead as bilateral introductions. **Step 5 (reproduction).** No agency gains by integrating unilaterally; recognised coordination problems are answered by adding a further coordinating body or a further proposal; and each programme builds a constituency and a reporting routine that raise the cost of later consolidation, the increasing returns associated with path dependence (Pierson, 2000). The configuration persists until an external mandate changes the pay-offs.

Status of the claim. This is an explanatory conjecture reached by inference to the best explanation, not a natural experiment: there is no exogenous shock and no counterfactual. The documents show the patterns the mechanism predicts; they do not show the agencies' pay-offs or decisions, and no formal game is specified. The brief does not claim that capital is irrelevant, that agencies are negligent, or that fragmentation is the binding constraint. The conjecture would be refuted if agencies evaluated on shared cohort outcomes remained fragmented, if cohort outcome data appeared in agency reporting without any integration of records, or if private intermediaries were found routing founders between public programmes with a record.

4. Method

Design. Qualitative document analysis (Bowen, 2009) within an abductive logic of theory construction (Timmermans & Tavory, 2012). Documents are treated as artefacts whose selection of metrics, omissions and inconsistencies are data about how their producers are organised and held to account (Prior, 2003).

Corpus and inclusion criteria. Sources were eligible if they (a) were produced by, or directly report the output of, a Saudi public or quasi-public body with an entrepreneurship, SME, venture-finance or innovation mandate, or an international body assessing the Saudi ecosystem; (b) were published 2019–2026 (government data files may cover earlier years); (c) report programmes, metrics, mandates or performance; and (d) were publicly accessible. The pilot corpus has three parts. First, eleven documents accessed 4–5 October 2026, of which only the Ministry of Investment’s 2026 startup handbook (read from a public mirror), the SME authority’s Q4 2024 monitor and one index page were read in full; the rest are official extracts or press reports of official figures. Second, the research authority’s January 2025 Markets of Tomorrow Accelerator report with the World Economic Forum (RDIA & WEF, 2025), read through a page-referenced extraction (§). Third, nineteen government data files, whose provenance is logged in a project register for only some (Section 7), read against the national Open Data Handbook (March 2023; publisher not stated in the file); a file is coded only if its title, unit and years are stated. Each source was assessed on Scott’s (1990) criteria of authenticity, credibility, representativeness and meaning. The author’s evidence register screens 193 documents from two research folders: 126 retained (19 coded corpus, 38 descriptive data, 46 theory, 10 comparators, 13 analogical support), 62 excluded and 5 method artefacts; content was verified from the text for 61, classified from title and known issuer for 69, and from title alone for 63 [internal]. This brief rests on the subset read in full. Register tiers map to the tags used here as follows: T1 and T3 official sources to [OI] (T3 outcomes are self-reported), T2 scholarship to [PR], T4 to [IP]; T5 sources are normally excluded.

Coding and tiering. A first cycle applied an a priori frame (metric type: input, activity, programme output, cohort outcome; source named; coordination claim; systemic routing; brokerage). A second, inductive cycle added numerical discordance and unattributable records. A discrepancy is coded only where quantity, unit and period are identical on the face of the sources; absence codes are applied only to sources read in full, because an extract cannot show that something is missing. Every source is tiered. Practitioner material, such as the StartupBlink 2026 index, a MENA early-stage data handbook and a German startup-association agenda, is used descriptively or for analogy only and carries no evidentiary weight here.¹

Hostile-witness logic. An agency has an interest in presenting its programmes as successful. Its documents are therefore most probative when they work against that interest: when they omit metrics that the agency’s own mandate implies it should hold, when they contradict themselves or a sibling agency on the same quantity, or when they concede a gap. These are treated as unintended disclosures about the measurement system, and inconsistencies are reported as findings rather than corrected. Genre is controlled: a promotional handbook is not expected to report failures, so absences carry more weight in statistical monitors and open datasets than in investor-facing material.

Positionality. The author is an ecosystem practitioner with a prior commercial interest in shared digital coordination infrastructure and holds angel investments in Saudi startups. A practitioner who has built coordination services has reason to find coordination deficits. Safeguards: every coded claim is tied to a page, dataset or dated source; disconfirming evidence was searched for and is reported in Section 5; rival explanations are stated in Section 6; no prescription is made. Platform operational data are not used as evidence; at most, anonymised aggregate context will describe how survey respondents are reached. Single-coder analysis remains the principal weakness. One core source, a private mapping of support programmes (Section 5.4), carries on the corpus copy a copyright notice naming the author; whether the author or an affiliate produced it is to be confirmed, and until then its independence is **UNSUPPORTED**.

5. Evidence

5.1 Capital abundance is documented; conversion is not

Every source that discusses venture finance reports abundance. The state venture vehicle reports annual venture investment rising from about USD 60 million in 2018 to USD 1.66 billion across 254 transactions in 2025 (SVC, 2026) [OI]; the investment ministry reports that Saudi Arabia captured more than 55 per cent of regional venture funding for three consecutive years (MISA, 2026, p. 8†) [OI]. The project’s open-data register also transcribes a series of announced local technology investment of SAR 2.30, 2.296, 2.10, 2.61 and 17.36 billion for 2021–2025 (Register, D01†) [OI], flagged there for a one-off check and because “announced” is not deployed. No source in the corpus reports survival, follow-on financing, revenue or employment for cohorts of supported firms, whether by programme, by agency or across agencies. The closest approximations are single-programme listings, such as thirteen growth-programme

companies listed (MISA, 2026, p. 18†) and ten offered on the parallel market in 2024 (Monsha’at, 2025, p. 34†) [OI], and one funnel. A vocational-training corporation’s entrepreneurship dataset (2019–2021, created September 2022) reports graduates (5,066; 4,282; 1,117), funded projects (4,889; 4,163; 1,973), projects under establishment (1,186; 1,221; 529) and projects opened (1,690; 1,632; 666) [OI]. The stage units differ (people and projects), funded projects exceed graduates in 2021, and nothing is recorded after opening, so it is a funnel of outputs, not a cohort. Conversion is measured from outside the agencies: WIPO ranks Saudi Arabia 31st on innovation inputs and 61st on outputs in 2025 (WIPO profile, quoted in the master paper, p. 2) [OI], and a practitioner dataset of 10,276 MENA startups computes conversion and failure rates by cohort and country from firm-level records (Clear, 2023, pp. 6, 15, 29) [IP], which shows that such measures can be built. The abundance is specific to equity: SME lending was 9.4 per cent of bank credit in 2024 and 10.8 per cent in the second quarter of 2025 against a 20 per cent aim (SAMA, via Arab News; Evidence Brief §4; master paper, p. 7) [OI].

5.2 Coordination mandates are stacked

The investment ministry lists “government coordination” as a service and describes market entry “through clear processes, digital systems, and direct institutional coordination” (MISA, 2026, p. 14†) [OI]. The research authority was given responsibility for coordinating research institutions (DIT, 2022†) [OI], and Monsha’at’s mandate places cross-sector collaboration at its centre (Monsha’at, n.d.†) [OI]. The same handbook presents ten “enablers”, each with its own programmes, eligibility rules and in several cases its own digital entry point, as a “comprehensive support system” (MISA, 2026, pp. 5, 14–27†) [OI], yet describes no mechanism by which a founder’s record or eligibility in one programme is visible to another. The pattern extends beyond startup support. A government co-creation register (file “Statistics of CoCreation II”; publisher not stated in the file) lists 280 hackathons, challenges, workshops, laboratories and initiatives from 2019 to 2026 run by 68 organisations, 40 of which ran at least one hackathon or challenge [OI], and the digital government authority’s innovation centre alone lists 19 services, including an idea bank, open-innovation sessions and a challenge programme (DGA services file) [OI]. A sector fund also runs its own voucher scheme of enabling services for cultural enterprises in seven fields, among them strategy, shared and private workspaces, costing and HR models, marketing, legal and compliance advice, financial audit and digital services (Cultural Development Fund, Namaa guide, slides 1–6†) [OI], which overlaps the SME authority’s consultation and mentoring offer. Coordination is distributed domain by domain; no body holds a mandate to integrate the integrators.

5.3 Measurement without a shared layer

Reported metrics are agency-level counts that cannot be summed, de-duplicated or followed through time. None of the nineteen data files contains a firm or founder identifier, and none reports firm survival or exit. This silence is not itself evidence of fragmentation. Open data are aggregated by design; the national handbook frames openness as licence, machine-readable format and discoverability, defines metadata only in its glossary, and says nothing on identifiers or record linkage (Open Data Handbook, 2023, pp. 4, 10–11, 15) [OI; tier provisional]; and it advises that “rather than trying to estimate which data will be most valuable, it can be helpful to look at which data will be easy to release” (p. 10), which predicts aggregate counts, although SDAIA’s companion toolkit (March 2023) adds a data value assessment step (SDAIA, Open Data Toolkit, slide 1†) [OI]. It also notes that openness must be balanced against personal privacy (p. 8). The finding on linkage therefore rests on the agencies’ own documents, which describe no route between programmes (Sections 5.2 and 5.4), not on the open data. The strongest hostile-witness evidence is discordance: the same quantity reported with different values across, and within, official channels (Table 1).

Table 1. Reported values for the same quantity in official channels

Quantity	Value A	Value B	Source and tier
Venture investment, 2025	USD 1.66 bn; 254 deals	USD 1.7 bn; 257 deals	SVC year-end release vs SVC impact report as reported in press [OI]
Same, within one document	USD 1.66 bn “committed” (p. 8)	USD 1.66 bn “deployed” (p. 12); chart shows USD 1.3 bn, period unlabelled (p. 12)	MISA 2026 handbook† [OI]
Cumulative VC since 2023	More than USD 3 bn (p. 8)	More than USD 3.4 bn to Q3 2025 (p. 12)	MISA 2026 handbook† [OI]

Quantity	Value A	Value B	Source and tier
Support-centre reach, 2024	41,076 “SMEs benefited”	46,204 beneficiaries, of whom 7,000 enterprise owners; 46,211 in a second dataset	Monsha’at 2025, p. 34† vs Open Data Platform† [OI]
Jadeer qualified enterprises, 2024	2,100	2,644	Monsha’at 2025, p. 34† vs Open Data Platform† [OI]
Academy trainees, 2024	51,110	45,507	Monsha’at 2025, p. 34† vs Open Data Platform† [OI]
NTDP launched startups	Defined as cumulative	Annual values 410, 463, 471 sum to the stated total of 1,344	MCIT open dataset† [OI]
NTDP entrepreneurs supported	Defined as cumulative since establishment to 2025	Annual values 1,023, 1,327, 1,206 sum to the stated total of 3,556	NTDP file in corpus; Register D03 [OI]
Innovation labs, 2022–2025	Title: “yearly”	Column header: “cumulative”; values 19, 24, 26, 28	Innovation-labs file in corpus [OI]
Established business ownership, 2024	19% (GEM official communication)	“nearly 14%” (companion launch release)	GEM, as recorded in master paper, p. 5 [OI]

None of these differences need be an error; committed and deployed capital, preliminary and final counts, and individuals versus enterprises are legitimate distinctions. The analytic point is that no shared layer reconciles them at source, and that the cumulative-or-annual definition is unstable even within one programme family.

Recorded, not coded. Three further observations meet the rule that quantity, unit and period must match, and are therefore not coded as discordance. A regional and a sectoral investment file total 2,363 and 1,730 and carry no title, unit or year (Register, D08, D09) [OI]. A 2025 service-provider file reports 47,657 sessions, while a 2025 beneficiary dataset reports 47,656 beneficiaries (†), under different stated units. The StartupBlink rank series 66, 65 and 38 for 2023–2025 appears in an open-data file (Register, D02; StartupBlink, 2026, p. 155) [IP] and must not be confused with the 66th-to-46th rise in the WIPO index (master paper, p. 2).

Unattributable records. Monsha’at’s 2025 service-provider file (281 rows) reports 47,657 sessions [OI]. Of these, 29.5 per cent (137 rows) are attributed to an “unidentified” centre and 8.8 per cent to an unidentified region; enterprise owners account for 19.2 per cent of sessions, people “interested in entrepreneurship” for 57.3 per cent, idea owners for 5.9 per cent and “other” for 17.6 per cent. Three variant labels appear to denote one business centre, and region names mix Arabic and English. The unit is the session, so the file cannot say how many enterprises were served. The co-creation register records whether a result exists, not what it was: of 265 completed activities, 230 are flagged as having a result and 35 not, overdue by 2 to 2,184 days (median 64) [OI].

The research authority as hostile witness. The Markets of Tomorrow Accelerator, run from 2023 across 36 workshops (RDIA & WEF, 2025, pp. 10, 27†) [OI], defines 22 enabling-environment indicators (pp. 22–25†), rates five as strong and five as weak, and leaves twelve unrated, including both finance indicators (pp. 23, 25–26†), yet describes the financial ecosystem as “prosperous and dynamic, with accessible innovation funds, venture capital” (pp. 29, 43†): capital adequacy is asserted, not scored. It states that 93 indicators were captured but never lists them and publishes no scores or weights (pp. 26–28†), says 32 market reports informed its sizing while its bibliography lists 16 (pp. 17, 45–46†), and draws 2030 charts whose bars contradict the accompanying text (p. 19†). Of 402 initiatives catalogued, 287 were public-sector; 61 of those (21 per cent) were relevant to the two target markets, against 72 of 83 private (87 per cent) and 28 of 32 academic initiatives (p. 27†). The selection rules are not stated, so this is a possible fragmentation signal, not a finding. The report covers two target markets, not startup support in general, and is used here only for what it shows about measurement.

5.4 Coordination is enacted bilaterally

Where documents describe how founders are connected to resources beyond a single programme, they describe brokerage: “curated introductions to VCs, giga-projects, and corporate innovation units” (MISA, 2026, p. 14†) [OI]; state-owned-enterprise programmes offering “connections” to their own units (pp. 20–22†) [OI]; an investor in Monsha’at’s monitor reporting more than 50 corporate–startup matchmaking activities in a year (Monsha’at, 2025, p. 28†) [OI]. Brokered connections are valuable but bilateral; the knowledge they generate stays with the parties. The

corpus contains no systemic routing, meaning a rule or infrastructure by which a founder declined or graduated in one programme is referred, with a record, to the next.

The structure measured from outside. The only ecosystem-level network measurement in the corpus is a private mapping of support programmes, compiled from secondary research (Annual Report on Entrepreneurship Support Programmes in the Kingdom, 2025, third edition, p. 7) [IP] [provenance to be confirmed, Section 4]. It records 418 programmes launched in 2025 (369 in 2024) by 68 owner entities, within a base of 144 distinct entities of which 93 appear as partners and 17 as both (pp. 7, 13). Partnership links doubled to 376 without the network widening: ten owners make 221 of them (58.8 per cent), and one makes 63 through six partners (p. 32). Only 135 of 6,324 possible pairs of entities are linked (2.13 per cent), the largest cluster holds 99 entities, and twelve clusters stand apart (p. 32); 137 programmes (32.8 per cent) have no registered partner, and 203 of the 281 with a partner have exactly one (p. 33). Only 53 of the 144 entities appear in both 2024 and 2025, although the report cautions that absence means no new programme was launched, not withdrawal (pp. 13–14). Of 321 programmes classified by entry stage, 253 (78.8 per cent) begin at the idea stage against 9 at fast growth (p. 24), and no programme extends across all four stages (p. 55). The report states that it measures the volume of support, not beneficiaries, funding, duration or impact (p. 18), and records supply, not demand (p. 57). It therefore quantifies bilateral structure and the missing outcome layer, but is itself a measurement made from outside the agencies, which is the point at issue.

5.5 Disconfirming evidence, reported in full

Several observations cut against the reading. Monsha’at’s two open datasets for 2023 and 2024 agree closely with each other (35,106 and 35,108 beneficiaries; 46,211 and 46,204; author’s coding matrix†) [OI], so the agency can reconcile its own statistics and the discordance arises between publication channels, not within the statistical function. The vocational-training dataset shows that one agency tracks beyond participation to project opening, so the absence of outcome data is not universal; the paper’s claim must be stated as the absence of any cross-agency cohort outcome. Private intermediation is thickening: venture investors in the market rose from 34 in 2018 to 200 in 2025 (Entrepreneur Middle East, 2026, reporting SVC) [OI]. Measured activity is growing fast: StartupBlink gives Saudi Arabia 97.0 per cent ecosystem growth, the highest in the G20, and Riyadh 117.6 per cent (StartupBlink, 2026, pp. 155–156) [IP], which a naive persistence reading would not predict. That index is partner-influenced: the national technology programme is a listed local partner (p. 12), the publisher says it has partnered with over 100 government agencies and that rank gains can come from newly integrated data (p. 429), and it states that it excludes surveys, interviews and causal models (p. 429); whether the partnership supplies data is not stated in the extract. Two official bodies also name the gap. The research authority’s report ranks a cross-sector observatory portal for data sharing and activity coordination fourth of eight short-term actions (RDIA & WEF, 2025, pp. 36, 39†) but assigns no owners, although it lists ownership as a guiding principle (p. 30†; master paper, p. 12). The National Strategy for Entrepreneurship and SMEs approved on 15 September 2026 includes an SME database “linked to other databases” and an impact-monitoring framework, and names a data gap (SPA, 2026b†) [OI]. These concessions support the measurement finding; delivery on shared identifiers and cohort outcomes would weaken the equilibrium. The strategy text itself was not published at the time of writing.

5.6 The UAE counter-case (MBRCGI 2015 framework)

The UAE adopted the full guidance-plus-mandate model: a 2015 Government Innovation Framework from the Mohammed bin Rashid Centre for Government Innovation (MBRCGI) and a Chief Innovation Officer (CIO) in every federal entity. A decade later the OECD records “multiple ministries conducting training programmes on AI or digital skills”, “policy silos” and “a fragmented promotional landscape” (OECD, 2026, chs. 6, 8; UAE Evidence Review, p. 5) [OI], and a peer-reviewed study ranks overlapping functions, poor coordination and lack of integration between government entities among critical barriers to service innovation (Bashir, 2025; Review, p. 6) [PR]. The framework’s second edition (18 November 2024) devotes a phase to “enhancing government coordination and integration” and formally “empowered” the CIO role (Review, pp. 3–4) [OI], an official concession that coordination remained unfinished. The UAE then moved to infrastructure with mandatory participation: a Unified Digital Platform Policy and a data-interoperability framework with “mandatory requirements” (Review, pp. 7–8) [OI].

Handled honestly, the case supports one half of the argument only. It is evidence that guidance and role mandates did not dissolve fragmentation. It is not evidence that infrastructure does: no evaluation of the pivot’s effect on startup outcomes was found. On StartupBlink’s 2026 country table the UAE ranks 30th (down nine places, growth of –27.5 per

cent) and Saudi Arabia 28th (up ten, +97.0 per cent) (StartupBlink, 2026, pp. 25, 155) [IP], but the index scores tax parameters within a business-environment component worth 13.33 per cent of the total (pp. 434–435) and responds to data integration (p. 429), so it cannot serve as an outcome evaluation of either country’s coordination infrastructure. An earlier exploratory comparison of business-environment and overall ranks is withdrawn from this brief because its underlying country table is not in the corpus. The claim that the CIO role became ceremonial is **UNSUPPORTED** (Review, pp. 4–5, 13). The OECD passages concern investment promotion and skills programmes, not startup support, and were captured from indexed snippets (Review, p. 14). One peer-reviewed source credits federal–emirate fragmentation with spurring competition among emirates (Review, p. 14), which the paper must engage.

6. Strongest objections and replies

Maturity lag (WIPO). *Objection:* outputs follow inputs with a delay, and the evidence of catch-up is already visible: the master paper records a rise from 66th to 46th in the WIPO index over 2021–2025 (p. 2) [OI], and StartupBlink reads the gap as outcomes “now catching up” (2026, p. 442) [IP]. *Reply:* the design cannot exclude lag, and the two explanations are not exclusive: an equilibrium would slow the closing of a lag rather than prevent it. They predict different time paths. Lag predicts that the gap narrows as cohorts age regardless of agency behaviour and that headline figures converge; the equilibrium predicts narrowing slower than input growth and continued discordance across editions. The master paper itself records a fifteen-place jump in a single year followed by five places in three years (p. 8) [OI], which reads as a plateau only if the index model is stable, and WIPO warns that year-on-year comparisons must account for model changes (master paper, p. 7). The argument does not rest on WIPO; the absence of cross-agency cohort outcome measures stands independently. The peer benchmark is partly withdrawn: Korea’s 4th-on-inputs, 6th-on-outputs split is recorded (master paper, p. 2) [OI] but compares a mature economy, and the claim that the UAE shows “no comparable divergence” rests on its composite rank, not its input–output split, and is **UNSUPPORTED**.

“Coordination will emerge naturally.” *Objection:* private investors multiplied from 34 to 200 and activity is growing fastest in the G20 (Section 5.5). *Reply:* emergence through a dense private intermediary layer is one of the stated boundary conditions under which the equilibrium weakens, so the objection names a falsification route rather than refuting the claim. Private brokers route bilaterally and keep the records. The deliberate-redundancy variant, that multiple entry points widen access and allow experimentation, predicts that redundant programmes are evaluated comparatively and the weaker closed, which requires the shared outcome layer whose absence is documented.

Guidance without infrastructure (UAE), and its mirror. *Objection:* if the UAE shows guidance fails, Saudi Arabia already has the infrastructure: a national single sign-on with more than 530 connected platforms and 23.5 million registered users (SDAIA, November 2024, as recorded in the Evidence Brief, §5) [OI] and a business centre whose registration automatically enrolls an establishment with several authorities (Evidence Brief, §5) [OI]. *Reply:* this forces a narrower claim. Identity and transaction infrastructure exists; what is missing is a shared firm and founder outcome layer across support agencies. Why the earlier business-centre unification did not generalise to support programmes is unresolved and is a gap in the current argument.

Cost of inaction. *Objection:* if the cost of fragmentation cannot be stated, there is no case for acting on it. *Reply:* the cost cannot currently be quantified from public sources. An earlier working-paper estimate, which applied the gap between the SME share of GDP and its 35 per cent target to national GDP and projected it to 2030 (master paper, p. 19), is withdrawn as **UNSUPPORTED** at peer-review standard; the UAE review found no sourced figure for duplication costs either (Review, pp. 14–15). The unquantifiability is itself predicted by the mechanism, since cohort outcomes are not observed. The planned survey measures a behavioural “fragmentation tax” (duplicated applications, re-entered data, reporting burden), which is a firm-level burden, not a macroeconomic cost. The paper will make no cost claim.

7. Limitations

The corpus is a pilot, predominantly English or aggregate, with several sources read only as extracts or press reports; the handbook was read from a public mirror and the research authority’s report through an extraction. Absence of outcome metrics may therefore reflect genre rather than governance, and the corpus cannot yet rule this out; the same holds for firm identifiers, whose absence from open data is expected and is not counted as evidence. Provenance logging is incomplete: the project’s open-data register covers 11 files, rating six as usable, two as held for missing title, unit and year, and three as excluded, and it logs no dataset for its claim that agencies duplicate or leave gaps between

one another (claim C4), so the mechanism rests on inference from absence and discordance, not on direct measurement. Analysis is single-coder by an author with a stated interest. Documents record what institutions say, not what they do; internal referral or data-sharing arrangements may exist unpublished. The equilibrium is inferred, not observed, and the design cannot estimate any effect on outcomes. Capital abundance is established for venture equity only, so the working title states the claim more strongly than the evidence supports, and even the research authority's report asserts rather than scores financial adequacy. Fragmentation is not unique to state-led systems: a German startup-association agenda describes "currently fragmented responsibilities" for skilled immigration (p. 4) and over 500 independent administrations deciding identical matters differently (p. 11; PDF pages) [IP], which concerns immigration, not startup support, but shows the pattern in a market-led federal state. The claim that the mechanism is stronger where the state is the dominant resource provider is therefore **UNSUPPORTED**, because no market-led comparator on agency measurement is in the corpus. Of the register's 193 documents, 132 are classified from title or title and issuer alone, and the Section Map in the register shows no source mapped to the master paper's introduction or recommendations, so the retained corpus is largely unread. The Monsha'at monitor cited in Table 1 is logged in the register as Q3 2024, not Q4 2024, and must be reconciled. Several literature references are not yet verified against source texts, the 2025 GEM figures come from a news-agency report, and the 2026 strategy text was unavailable. The UAE evidence review relies partly on indexed snippets.

The survey data are not yet fielded (Appendix A). The seven-role item bank is a Tranche 1 draft (3 October 2026; 48 items, 43 for live fielding); founder items and role blocks are not drafted, and its planning estimate of about 18 minutes exceeds the 15-minute ceiling it sets for senior respondents (Item Bank README) [internal]. None of the survey hypotheses, including the binding-constraint test (H7), has been tested. The draft decision rule for H7 exposes a design flaw: it names finance and human capital as the refuting domains and is silent when markets, policy, culture or supports outrank coordination, although the item bank offers all of them as barrier domains (LIST_DOMAIN). A low rank for finance is what the thesis already predicts, so the test lies in coordination exceeding every rival domain, markets included. The rule will be rewritten to cover every rival domain before fielding.

8. Questions for the reviewing professor

Ordered by how much the answer could change the paper.

- 1 **Genre and timing.** Is a theory-building document analysis on a pilot corpus publishable in ERD, or should I hold submission until a sampled corpus of 30 to 40 documents, including Arabic originals, has been double-coded, or until the survey is fielded with a corrected H7 decision rule?
- 2 **Novelty.** Is the fragmentation equilibrium a contribution to ecosystem theory, or a relabelling of whole-of-government coordination failure and of the territorial-approach literature, or of the digital-infrastructure strand of Sussan and Acs? Which literature would you expect a referee to say I have missed?
- 3 **The claim and the title.** Should I abandon "Capital Was Never the Constraint" for a conversion and governance claim, given that abundance holds for venture equity and not for SME credit, and that even the research authority asserts rather than scores financial adequacy?
- 4 **Hostile-witness evidence.** Will a referee accept omissions, unattributable records and cross-channel discordance in official documents as evidence about governance, or dismiss them as genre and definitional artefacts, given that one agency does track a funnel to project opening and that open data are aggregated by design? Is "no cross-agency cohort outcome" a defensible form of the absence claim, and what additional test would persuade you?
- 5 **The UAE case.** Given that it shows the insufficiency of guidance but not the sufficiency of infrastructure, should it remain as a counter-case, be cut, or be replaced by a documented failure of a national coordination layer?

Appendix A. Primary survey: status, questionnaire by stakeholder and analysis plan

A.1 Status of the primary data. No survey response is held in the project corpus and nothing has been analysed. What exists is the instrument: the GCC Ecosystem Research Survey v3, Tranche 1 item bank (draft 1, 3 October 2026; English

only; for pilot), with sheets for flow and rules, role tokens and vignette, items, lists and scales, and an analysis plan [internal]. It holds 48 items: 43 for live fielding (P1) and 5 pilot reserves (P2). An earlier six-role live form exists; its responses are not in the folder and are not used. Not yet drafted: the role blocks (role-specific mirror items and count bands), founder items F1–F8, Modules 4–7, and the shared closing carried over from the earlier version (adoption and willingness-to-pay items). The planning estimate is 17.7 minutes for the live path, of which 10.6 minutes are drafted items and 7.2 minutes are placeholders for the undrafted blocks, against a rule-of-thumb ceiling of about 15 minutes for senior respondents; none of these is a measured time. The item bank sets no target sample size or sampling frame, only reporting thresholds: a track-level result needs *n* of at least 30, smaller tracks are “indicative only”, and cells with *n* below 5 are suppressed.

A.2 Design. Seven roles enter through one shared spine; stems are identical and only the noun phrase changes (Table A1), so cross-role differences reflect role position plus the referent. The order is fixed to avoid priming. After routing and country, a 100-point allocation across seven barrier domains (policy, finance, culture, supports, human capital, markets, coordination) and a forced-choice priority item precede any item using the words fragmentation, duplication, coordination, platform or verified. A 2×2 vignette follows (warm or cold introduction by verified or self-reported record; two cells per respondent; likelihood of a next step from 0 to 10), then the role matrix (organisations worked with by type, how new relationships began across ten chips, an access wish list and its main obstacle) and the decision block, a mirror of founder and decider experience gated by D1. The fragmentation-tax block asks about organisations exchanged paperwork with, personal hours per month, the share re-entered and the expected saving if reused. A Likert core covers duplication, visibility, verification, merit versus network and standardisation, with a forward and a reverse item per construct and one attention check; then willingness to contribute six types of data and the conditions, the undrafted role block and closing, and covariates (including how the respondent found the survey and closeness to the inviter) with an unlinked follow-up opt-in. “Don’t know”, “not applicable” and “prefer not to say” are never midpoints. Invitations are one-time links carrying role and wave only, with no name, employer or contact field.

Table A1. The questionnaire by stakeholder (shared spine as worded for each role; role-specific blocks pending). Source: item bank sheets [internal]; role-specific item wording is not in the folder, only earlier-version identifiers.

Stakeholder (organisation types offered at routing)	Shared items filled with: {ACTOR} · {COUNTERPART} · vignette next step · standardisable item	Decision block (D1–D5)	Role items not yet drafted (earlier-version IDs)
Founders and startups (6 stage types: idea; product built, pre-revenue; early revenue; scaling; bootstrapped and profitable; other)	“my startup” · “the programs, investors, mentors and corporate partners I deal with” · “agree to a first meeting or submit an application” · “my application and company information”	Asked whether turned down; counts submissions; ranks perceived reasons; share told the reason	F1–F8, including F2 (visibility), F4 (standardisation), F6a–b (paperwork count and time)
Incubators and accelerators (9 types, from pre-incubation to corporate and government-run programmes)	“my organization” · “the startups and founders I evaluate or support” · “invite them to an interview” · “applications, intake and reporting”	Asked whether declined or passed; counts reviewed (large-volume band); ranks reasons given; share told	C2, C5, N1, N2
Investors (9 types, from seed VC and angels to family offices, CVC and and sovereign-backed)	“my firm or investing activity” · “the startups and founders I evaluate or invest in” · “take a first meeting” · “due-diligence information (for example a standard data room)”	As incubators	Q2b, C5, N1, N2
Government and regulators (9 types, from central ministry and SME authority to free-zone authority)	“my agency” · “the startups and programs my agency funds, regulates or supports” · “add them to a funding or support shortlist” · “program reporting and performance indicators”	Same wording; “not applicable” expected (rule RL2)	Q2, C5, N1
Universities and academia (9 types, from entrepreneurship centres and TTOs to research institutes)	“my institution” · “the startups, industry partners and investors my institution works with” · “agree to a first partnership or licensing discussion” · “spinout and IP licensing terms”	Same wording; “not applicable” expected (RL2)	C2, C5
Mentors and service providers (10 types, from operators and lawyers to fundraising advisers)	“my practice or firm” · “the founders and startups I advise” · “agree to a first advisory session” · “founder intake and engagement terms”	As incubators	C2

Stakeholder (organisation types offered at routing)	Shared items filled with: {ACTOR} · {COUNTERPART} · vignette next step · standardisable item	Decision block (D1–D5)	Role items not yet drafted (earlier-version IDs)
Corporates and industry partners (10 sectors, from banks and telecoms to public-sector enterprises)	“my company” · “the startups my company evaluates or pilots with” · “agree to scope a pilot” · “pilot agreements and procurement documents”	As incubators	C5

A.3 Constructs, tests and pilot. H7 (binding constraint) uses the allocation, the priority item and the reason ranking: within-respondent Wilcoxon signed-rank tests of the coordination share against each other domain (six contrasts, Holm-adjusted), Kruskal–Wallis by role and ordinal logistic regression. H1 (visibility and verification) uses the visibility and verification items, the main-obstacle item, the vignette verification premium and the reason “information could not be verified”. H3 (merit versus network) uses the gap between perceived and preferred merit-based access, a cost-bearing trade-off, the share of relationships begun through personal introductions and the vignette introduction premium, estimated with a cumulative-link mixed model (random intercept per respondent). H2 (duplication, standardisation, fragmentation tax) uses duplicated hours, personal hours times the share re-entered, with bounds from band limits, in hours not currency. H4 (access gaps) uses the role-by-type tie matrix and wish list, and compares founders with deciders on the mirror items (Mann–Whitney, Cliff’s delta). H5 and H6 (data contribution) use willingness across six data types (Friedman test), the conditions (independent governance, reciprocity, mandated participation) and a behavioural anchor. Global rules: Holm correction within hypothesis families, medians and distributions for ordinal data, effect sizes with 95% intervals, and a time-stamped pre-registration before fielding. A pilot of about 30 respondents swaps an item for its reserve if its item-rest Spearman is below .30 or over 80% of answers fall in one category; five think-aloud interviews per role test the tokens; Arabic follows by forward and back translation after wording lock [internal].

A.4 What the instrument cannot yet test. No drafted item asks agencies how they are evaluated or what they would lose by integrating, so Steps 1, 2 and 5 of the mechanism in Section 3 are untested; Step 3 is touched only by willingness to contribute and reciprocity, and Step 4 by the visibility, verification and vignette items. The Government role block is where agency-accountability items belong. Recruitment is by invitation, so network-recruited bias is testable through the two recruitment covariates but cannot be removed, and no weighting or population inference is planned. Governments and universities will often answer “not applicable” to the decision block, and the founder branch of that block is selected on having been turned down. Willingness and time are self-reported, hypothetical willingness overstates behaviour, and the top bands rest on assumed caps. The form is English only until translation, which repeats the English-language bias of the document corpus. Three items remain unreconciled with the earlier version: the “is” wording of the merit item, the conditions item, and the vignette draw, which needs custom scripting (fallback: one vignette per respondent, pooled analysis only). The decision rule for H7 must still be rewritten to cover every rival domain (Section 7).

¹ Industry-practitioner sources are footnoted as such and used for analogy or description only: StartupBlink (2026), Clear (2023) and the German Innovation Agenda 2030. No consultancy survey is relied on as an empirical anchor in this brief.

Source register (corpus status). In the project folder and cited at page or dataset level: Stam & van de Ven (2021), *Small Business Economics* 56, 809–832 [PR]; Scaringella & Radziwon (2018), *Technological Forecasting & Social Change*, p. 1 [PR]; UAE Counter-Case Evidence Review (project synthesis; tiers given for its OECD 2026, Bashir 2025 and UAE official sources); StartupBlink GSEI 2026 Saudi extract [IP]; RDIA and WEF (2025) as extracted in the project’s page-referenced dataset (‡); nineteen government data files, the Open Data Source Register, the Open Data Handbook (March 2023; tier provisional) and the WP-2026-01 Evidence Register v1–4 (internal); read from the author’s Drive copies (†): the 2025 programmes report [IP], the Cultural Development Fund Namaa guide and the SDAIA Open Data Toolkit; Clear (2023) *MENA Early Stage Data Handbook* [IP]; *Innovation Agenda 2030* (Germany) [IP]; Survey v3 Item Bank, Tranche 1 (internal); WP-2026-01 master paper (24-page version), used only as a locator for figures it records; Innovation Clusters Evidence Brief. Not in the folder (†): MISA *Saudi Startup Ecosystem Handbook 2026*; Monsha’at *SME Monitor Q4 2024* (the register logs Q3 2024); SVC releases; WIPO GII 2024–2025 profiles; GEM releases; SPA (2026b); DIT (2022); the 2024–2025 Monsha’at and MCIT datasets read earlier; the original RDIA and WEF report; and, not read in full, the 2023 programmes report, the MISA Investor Guide, the Social Development Bank workspace guide and the 132 register rows classified without reading their content. The earlier ERD manuscript draft and revision tracker are no longer in the folder, so figures taken from them carry †. Reviewed and not used: the deep-tech report, the MITEF summary and the platform-theory papers; platform operational data are excluded. Literature cited without page (Isenberg 2010; Spigel 2017; Stam 2015; Sussan & Acs 2017; Olson 1965; Pierson 2000; Bouckaert et al. 2010; Christensen & Lægheid 2007; Bowen 2009; Prior 2003; Scott 1990; Timmermans & Tavory 2012) is characterised from the working bibliography and must be verified against source texts before submission.

Latest version (v4): <https://drive.google.com/drive/folders/1OrXrB3RVnNvXY9JqWZyxbfXqDc6FVGPC>