

Site logic and city persistence

A registered claim, refused rather than tested — and the corpus diagnostics that made the refusal necessary.

MEASURE CLASS

Class 3 — population claim

REGISTERED

2026-09-09

STATUS

Refused. Figures pending source verification.

FALSIFIER

§7 — what would refute or unblock this

ABSTRACT

Cities are not placed at random. A city usually sits where some physical fact of the ground — a river crossing, a harbour, a spring in a desert, a pass through mountains — made that spot necessary rather than the one fifty kilometres away. The question registered here is whether that founding reason, which this paper calls the city's **site logic**, predicts how long the city goes on existing. It is an attractive question and this paper does not answer it. It refuses to.

The refusal has two causes and both are structural rather than temporary. First, the only body of evidence available — a table of 59 Asian and Middle Eastern cities — has no column recording when or whether any city ended, so persistence is not a measurable quantity in it at all; the sole computable quantity is age since founding among cities a compiler in 2026 chose to profile, which counts survivors and never counts the dead. Second, the claim being tested was written from that same table, names four of its rows as its evidence, and appears restated inside the table's own descriptive fields. Testing it there would return the answer that was put in.

What this paper publishes instead is a measurement of the corpus rather than of the world: 16 site-logic categories declared against 15 actually used, 8 of those 15 holding two cities or fewer, all 59 rows resting on one unverified source, and a published summary of the set that omits precisely the categories which would have stopped the analysis. Those counts are real and recomputable. They make no claim about cities.

1 Background

Why any particular city is where it is has an unusually satisfying class of answer. Istanbul sits on a strait, Kabul in a pass, Palmyra on a spring in a desert, Kaifeng at the end of a canal. Each of these is a constraint the ground imposed on movement, and the city is the place where movement had to stop, transfer, wait or pay.

From that observation a stronger claim follows very naturally, and it is the claim registered in §2: if a physical constraint is what made a city necessary, then when the constraint disappears the city should decline. Steam removed the need to wait for a monsoon; the caravan route moved away from the spring; the river silted. On this account the city is a machine for solving a problem, and it fails when the problem is solved elsewhere. The claim usually carries an exception attached — that cities founded on sacred sites, which have no economic reason to

exist at all, are the durable ones precisely because pilgrimage does not depend on trade routes.

The strongest counter-argument, which should be stated before any data appears, is that the claim mistakes an effect for a cause. Once a city exists at scale it acquires reasons to persist that have nothing to do with why it was founded: a population, a capital's administration, roads and rails built to reach it, sunk capital that would be expensive to abandon. On this reading founding logic explains the *first* few centuries and almost nothing after, and any apparent relationship between site logic and longevity is really a relationship between site logic and how early the city was founded — old cities are old because they started early, not because their site kept working.

This paper cannot adjudicate between those two accounts, and the reason it cannot is the finding. The evidence available to test them is evidence assembled by someone who already believed the first one.

2 Claim, class and preregistration

REGISTERED CLAIM

Among cities admitted by a stated rule, the physical site logic under which a city was founded predicts how long the city persists — with sacred sites persisting longest despite the weakest economic case for their existence.

MEASURE CLASS

Class 3. The estate grades a quantitative statement by what it is a statement *about*: class 1 describes a named artefact such as a table or a file; class 2 compares a handful of named cases and claims nothing beyond them; class 3 asserts something about a population, which requires a written rule saying which members of that population were eligible to be counted. The claim above is about cities in general. It is therefore class 3, and a class 3 claim without such a rule cannot be published as a finding.

DECISION RULE, STATED BEFORE COUNTING

The study proceeds as a finding only if all four of the following hold. If any one fails, the study is refused, and the refusal is published with the corpus diagnostics that established it.

D1A written **admission rule** — numbered criteria stating which places were eligible for the set and which were excluded — exists for the city table, and was written before any count was taken.FAIL

D2A persistence variable exists as a recorded column with a stated basis, so that "how long a city lasted" is read from the data rather than judged from prose after the answer is visible.FAIL

D3The corpus carries at least two **extractionally independent** sources — two genuinely separate routes to the evidence, not one route travelled twice.FAIL

D4The claim under test was registered independently of the corpus: it does not cite the corpus's own rows as its evidence.FAIL

Outcome: **0 of 4**. The decision rule was written before the counts in §4 were taken, but not before the corpus was built — the corpus predates this paper by hours, and that ordering is treated as a threat to validity in §6 rather than waved through.

SCALE TYPE OF EVERY VARIABLE

Site logic is *nominal*: sixteen named categories with no order and no distance between them, so no arithmetic of any kind is performed on it. Founding year is *interval*, not ratio — there is no true zero on a calendar, so years may be subtracted but not divided. Age since founding, derived by subtracting founding year from 2026, is ratio but is only as good as the interval variable underneath it. Category membership counts are *absolute*. Confidence tier is *ordinal*. Nothing in this paper multiplies, averages or fits anything across these types.

3 Admission rule

There is none, and its absence is the first finding rather than a gap to be apologised for.

The estate's own lane plan for this material, `ops/ASIA-LANE-PLAN.md`, has a section headed "Admission rules — written before any count". It contains two: `ADMIT-ASIA-STATUS-001`, governing which status-holdings may enter, and `ADMIT-ASIA-OFFICE-002`, governing which offices may. Both are written to a good standard — each states what is admitted, what is refused, and what would show the rule itself to be wrong. Neither concerns cities. A search of the estate's markdown for the string `ADMIT-ASIA` returns exactly those two matches. Nothing governs the city set.

How the 59 cities were actually found can therefore be stated only in plain words: they were recalled and typed. That is what the database itself records. Every one of the 59 rows carries the same `source_ref` string — `korea-and-asian-cities.xlsx scaffold 2026-09-09 — UNSOURCED, from recall` — and the same confidence value, **T0**, which in this estate's grading means asserted and not yet checked against anything.

There is direct evidence in the data that no rule was operating, and it is worth stating precisely because it does not depend on anybody's memory of how the set was made. The typology table lists worked exemplars for each category — 47 named tokens across 16 categories, with no category left blank. Ten of those named places have no row in the city table at all: Khartoum, Muscat, Aden, Calicut, Peshawar, Islamabad, Baku, Dhahran, Fairbanks, and a group written as "Suez towns". The typology's own illustrations include cities the set excludes, and no criterion anywhere states why they are illustrations but not members. Three of the ten — Khartoum, Baku and Fairbanks — are not in Asia or the Middle East, so the set's geographic boundary is unstated as well as its selection rule. A rule that admitted Fairbanks as an example of resource siting and then declined to admit it as a city would be a strange rule; the more economical reading is that there was no rule.

4 Data

Table 1 is the data of this paper. Every row in it is a **class 1** quantity: a count over a named file on a named date, describing that file and nothing beyond it. The distinction matters, and it is the reason this paper has anything to publish at all. The contents of the corpus are T0 — unverified assertions about the world. Counts of that corpus are T1 — measured once, from a named artefact, by a stated command, and recomputable by anyone with the file. A statement about how many categories a table declares is true whatever is true of cities.

Table 1. *Corpus diagnostics. All counts taken 2026-09-09 from `ops/asia/asia.db` opened read-only, from `ops/truth.db`, and from `Downloads/korea-and-asian-cities.xlsx` read with `openpyxl`. Confidence **T1** means measured once from the named artefact by the basis given; **T0** means the underlying content is asserted and unverified. Where the sources disagree the band is published rather than a point.*

QUANTITY	VALUE	BASIS — WHAT WAS COUNTED, AND HOW	CONFIDENCE
Site-logic categories declared	16	<code>select count(*) from site_logic</code>	T1 measured
Categories used by at least one city	15	<code>select count(distinct site_logic) from city</code>	T1 measured
Declared and never applied	1	Left join <code>site_logic</code> to <code>city</code> ; the unfed value is <code>resource</code> , mechanism "Mineral or fuel extraction"	T1 measured
Used but not declared (broken references)	0	Anti-join of <code>city.site_logic</code> against <code>site_logic.logic_key</code>	T1 measured

QUANTITY	VALUE	BASIS — WHAT WAS COUNTED, AND HOW	CONFIDENCE
City rows in the set	59–60	Band, because the one artefact disagrees with itself: the spreadsheet holds 59 data rows and its summary sheet says 59, while its read-me prose says "60 profiles"	T1 measured
Rows at confidence T0	59 of 59	<code>select confidence, count(*) from city group by 1</code> — returns one group	T1 measured
Distinct source references across the set	1	<code>select count(distinct source_ref) from city</code>	T1 measured
Categories holding exactly one city	5	Group by <code>site_logic</code> having count = 1: frontier-capital (Beijing), hydraulic (Angkor), monsoon-wait (Malacca), pass (Kabul), portage (Palembang)	T1 measured
Categories holding exactly two cities	3	Same query, count = 2: canal-terminus, colonial-foundation, strait	T1 measured
Cities in a category of two or fewer	11 of 59	Sum of the eight categories above; 8 of the 15 used categories admit no internal comparison	T1 measured
Countries represented	26	<code>select count(distinct country) from city</code>	T1 measured
Largest single-country share	8 of 59 13.6%	China, by <code>group by country order by count desc</code> ; then India 5, Korea 5, Japan 4	T1 measured
Countries contributing exactly one city	12	Group by country having count = 1. A prior working note recorded 15; recomputed today it is 12, and the note is superseded rather than deleted	T1 measured
Founding year populated	59 of 59	Non-null count on <code>founded_or_first_attested</code>	T1 measured / T0 content
Founding years before year 0	25	Cast to integer, value < 0. The column is stored as text, so an uncast comparison silently misreads it	T1 measured / T0 content
Founding years after 1500	16	Cast to integer, value > 1500	T1 measured / T0 content
Columns on <code>city</code> recording an end, an abandonment, a continuity state or a population	0	<code>pragma table_info(city)</code> — the eleven columns are key, name, country, founding year, site logic, why here, water source, principal function, transformation prose, source ref, confidence	T1 structural
Sibling dated tables in the same database that <i>do</i> carry an end year	4	<code>pragma table_info</code> over each table: <code>dynasty</code> , <code>office</code> , <code>succession</code> and <code>tradition</code> all carry <code>to_year</code> ; <code>city</code> carries no equivalent	T1 structural
Declared categories absent from the workbook's published summary block	8 of 16	The "Cities by site logic" block on the summary sheet lists 8 categories; 7 of the 8 omitted have members, the eighth is the empty resource	T1 measured
Cities unrepresented in that block	9 of 59	The block's eight listed categories total 50 cities; $59 - 50 = 9$, and the omitted nine include all five singleton categories	T1 measured

QUANTITY	VALUE	BASIS — WHAT WAS COUNTED, AND HOW	CONFIDENCE
Exemplar places named by the typology with no row in the set	10	47 example tokens split from <code>site_logic.examples</code> and matched against <code>city.name</code> ; 12 fail to match, of which "Calcutta" resolves to the row Kolkata and "upriver marts generally" is not a place	T1 measured
Of those, outside the stated Asia and Middle East frame	3	Khartoum, Baku, Fairbanks — read from the same token list	T1 measured
Typology rows with no exemplars at all	0	Trim-and-empty test on <code>site_logic.examples</code> ; the absence runs one way only	T1 measured
Written admission rules governing the city set	0	Search for <code>ADMIT-ASIA</code> across estate markdown returns 2 matches, both in <code>ASIA-LANE-PLAN.md</code> , for status-holdings and offices	T1 measured
Instances named by the claim under test that are rows in this set	4 of 4	<code>universal_truth #25</code> in <code>truth.db</code> names Palmyra, Malacca, Kolkata and Jerusalem; all four resolve to city rows	T1 measured




Figure 1. A census of the vocabulary, not a result. Each bar is one declared site-logic category and its width is the number of cities carrying it. Outlined bars are the eight categories the workbook's own published summary block leaves out — every category with two members or fewer, nine cities in total, including all five singletons. **No trend line is drawn and nothing is plotted against time.** A fitted line asserts a relationship in a population, which is a class 3 statement, and §2 records that this study fails all four conditions for making one. There is also nothing honest to plot it against: the only candidate second variable is age since founding, which is the quantity §5 shows to be circular here.

5 Results

The decision rule returns **refuse**, at 0 of 4 conditions met. Each of the four failures is set out below, because a refusal that does not show its working is indistinguishable from a shrug.

D1 — no admission rule. Established in §3. Without one, every proportion computed over the set is a description of what a compiler found interesting. "Nine of 59 cities are harbours" is a true statement about a spreadsheet and an empty one about Asia.

D2 — persistence is not operationalised. The `city` table has eleven columns and none of them records an ending. There is no end year, no abandonment flag, no continuity state, no population. This is not an oversight peculiar to a young table: four sibling tables in the same database — `dynasty`, `office`, `succession` and `tradition` — all carry a `to_year`. Cities alone are recorded as though they only begin. The consequence is that "persistence" would have to be coded from the free-prose `transformation` field, after the answer is already visible, by the person who wrote that prose. That is not measurement.

D3 — one source, wearing 59 coats. All 59 rows carry the identical `source_ref` and the identical confidence tier. Fifty-nine rows from one recall pass are one witness, not fifty-nine. The estate distinguishes **extractional independence** — two genuinely different routes to the evidence — from a single route travelled repeatedly, and this is the second. A further recall pass by the same author would not repair it, and neither would a second spreadsheet copying the first.

D4 — the claim and the corpus share an author, and the claim cites the corpus. This is the failure that makes the refusal structural rather than merely premature. The claim under test is recorded in `truth.db` as `universal_truth #25`, at tier T2, reading: "Cities sit where a physical constraint made them necessary, and survive its loss badly — except sacred sites, which survive everything." It names four instances. All four — Palmyra, Malacca, Kolkata, Jerusalem — are rows in the `city` table. Their `transformation` fields restate the claim in the corpus's own voice: Palmyra, "Existed only because of one spring, and ended when the route moved"; Malacca, "Superseded once steam removed the need to wait"; Jerusalem, "Weak economic logic; the most durable city in this table". The claim's own mechanism text says Jerusalem outlasts "nearly every city in the table" — the claim refers to the table by name. Adding rows to the same scaffold does not fix this, because the fix required is a different corpus, not a bigger one.

THE RESULT WITHHELD, SHOWN SO THE READER CAN SEE WHAT WAS AVOIDED

Table 2 exists to make the refusal legible rather than to be believed. It is the answer the naive analysis returns, and it is put on the page precisely because it looks convincing.

Table 2. *The withheld result. Age is 2026 minus the recorded founding year — the only persistence-shaped quantity the table can produce. The median is the middle value of the sorted ages, chosen over the mean because with three-member groups a mean is dominated by one row. These figures are reported to be rejected, not used.*

GROUP	CITIES	MEDIAN AGE, YEARS	BASIS	CONFIDENCE
Whole set	59	1,264	Median of 2026 – founding year, all rows	T0 content
sacred	3	3,226	Jerusalem, Varanasi, Mecca	T0 content
harbour	9	455	Median within the category	T0 content
planned-capital	5	271	Median within the category	T0 content

Read without discipline, that is a striking result: sacred cities show a median age about two and a half times the whole set's and about seven times that of planned capitals. It reproduces claim #25 exactly. It is worthless, and the reason is that it is the same statement twice — the claim was written from these rows, so the rows agree with it by construction.

THE ROW THAT FITS WORST: SACRED, N = 3

Every honest study has a case that fits badly, and here it is not a case but the whole load-bearing category. **Sacred holds three cities: Jerusalem, Varanasi and Mecca.** Four things are true of it at once, and together they are the paper.

It carries the entire finding — the exception clause in claim #25 is the interesting half of the claim, and this is the only category that tests it. It produces the largest effect in the withheld run, a median age of 3,226 years against 1,264 for the set. It has three members, which is below any threshold at which a comparison means anything. And all three are among the most famous continuously inhabited pilgrimage cities on earth. There is no failed sacred city anywhere in the set. A compiler would have had to go looking for one, and nothing in the record suggests anyone did.

The category cannot produce a negative result. Its own definition forecloses one: the typology records the `the_catch` field for `sacred` as "Weak economic logic, overwhelming durability" — the conclusion written into the category label, before any city is assigned to it. A category defined by its outcome will be confirmed by every member it admits. That is the sharpest single artefact in this paper, and it is a property of the vocabulary, not of the world.

WHAT WAS MEASURED, SINCE IT WAS NOT CITIES

The obvious hostile question is what a refusal actually measured, if it declined to measure the thing it registered. The answer is the corpus, and the measurements are in Table 1: 16 categories declared against 15 used; one declared and never applied; 8 of the 15 used categories holding two cities or fewer and covering 11 of 59 cities; 26 countries with the largest contributing 13.6%; 59 rows at one confidence tier from one source; ten exemplar cities named by the typology and absent from the set, three of them outside its unstated geographic frame; and a published summary that lists 8 of 16 categories and drops 9 cities, including every singleton.

That last one deserves stating plainly rather than as a bullet. **The summary block omits exactly the rows that would have stopped the analysis.** A reader who consulted only the workbook's own summary would see eight well-populated categories, would have no reason to suspect a singleton problem, and would proceed. The workbook's read-me compounds it twice more: its prose says "Sixteen answers cover the whole set", which is false — fifteen do, and the sixteenth covers nothing; and it says the set holds "60 profiles" where the sheet holds 59. A single artefact disagreeing with itself in the direction of overstatement is the only band this paper can publish, because a band normally requires two sources and here there is one.

Those statements are class 1. They describe named files and are true whatever cities turn out to be like. They are also the contribution: they give `universal_truth` #13 — "Selection is the default condition of any historical corpus; a trend over an unruly corpus measures the compiler" — its first fully worked instance, and they move #25 from a quietly held T2 to a claim explicitly registered as untestable in place.

6 Threats to validity

1. **The decision rule was written after the corpus existed.** §2's four conditions were fixed before the counts in Table 1 were taken, but the corpus was built earlier the same day and its shape was already broadly known. A rule written with any knowledge of the material can be shaped, however unconsciously, to produce the verdict the author expected. The verdict here is a refusal, which is the direction that costs the author something rather than the direction that flatters them — but that is a mitigation, not a defence.
2. **The replication of 16 / 15 / 5 is not independent.** The loader that built this database, `ops/asia/load_cities.py`, states in its own docstring that 16 categories are declared and 15 used, names `resource` as the unfed value, and names the same five singletons. This paper re-derived those numbers from the database with SQL rather than reading them off the docstring. That is one extraction path measured twice, not two paths agreeing. Genuine replication would require a second reader building the counts from the spreadsheet without seeing either the loader or this paper.
3. **Selection, which is the subject and therefore inescapable.** The set contains Palmyra and Petra, both famous ruins. It contains no obscure dead city at all. Any persistence figure computed here has a **survivorship** denominator — it counts what lasted, or what is famous for not lasting, and never counts the ordinary settlements that failed quietly and left no name for a compiler to recall. This is the mechanism claim #13 describes, operating on the very corpus assembled to illustrate it.
4. **Founding dates are interval and contested, and treated here as neither.** "Founded or first attested" merges two different events: a deliberate foundation with a date, and the earliest surviving mention of a place that already existed. Ages derived by subtraction inherit that

merge. No source is named for any of the 59 values, so none can be checked.

5. **The stored type of the founding column is text, not integer.** An uncast comparison against a number does not fail loudly — it returns a plausible wrong count. Two of Table 1's rows were initially wrong for this reason and were corrected by casting; the correction is recorded here rather than silently applied, because the failure mode is invisible in the result.
6. **A prior working note gave a figure this paper contradicts.** A note carried into this study recorded 15 countries contributing exactly one city. Recomputed today the value is 12. The note is superseded, not deleted, and its existence is a reminder that a figure repeated between documents is not a figure verified between documents.
7. **The catalogue understates by an unknown amount.** Every defect named in this paper is one that was caught. There is no way to estimate from inside how many were not, and a refusal published to a high standard is more likely to be read as complete than a finding is.

7 What would refute this, and what would unblock it

A refusal is a claim and must be refutable like any other. What is registered here is that the claim in §2 cannot be tested on this corpus. That is wrong if any of the following is shown.

- A written admission rule for the city set, dated before the corpus was assembled, is produced. The refusal rests on its absence and a document would end it.
- An end-year or terminal-status column with a stated basis is shown to exist on `city`, or is derivable from a named source that is not the transformation prose. Persistence would then be a variable rather than a judgement.
- The four instances named by `universal_truth #25` are shown to have been registered before the city table was built, with a dated record. The circularity in D4 would then be apparent rather than real.
- An independent reader, working from the spreadsheet without sight of the loader or this paper, returns different counts for the 16 declared, 15 used, 5 singleton figures. That would refute the diagnostics themselves, which are the only positive content here.

WHAT WOULD UNBLOCK THE ORIGINAL QUESTION

Five things are missing. Each is named so that someone can supply one rather than argue about the whole.

1. **An admission rule for the city set** — call it `ADMIT-ASIA-CITY-003`; it does not exist. It must state the geographic frame, since the set already leaks to Khartoum, Baku and Fairbanks in its own exemplars; a date floor; and what makes a place a city rather than a settlement. Until it exists every category share describes the compiler.
2. **A persistence variable that exists as data.** An end year, or a status column with a declared basis — continuously inhabited, reduced to ruin, refounded elsewhere — and a rule stating which of those counts as failure, written before any city is coded.
3. **The dead cities.** The set holds no failed city that is not also famous. A persistence rate needs a source-defined population with its own admission rule — a gazetteer, an archaeological register — from which survivors and failures are drawn by the same criterion.
4. **A second, extractionally independent source.** Not a second recall pass and not a second scaffold copying the first. Until one exists nothing in the corpus rises above T0 and this paper must not imply that it does.
5. **An independent test set for claim #25.** A city set assembled under a rule written before #25 was formulated, containing none of Palmyra, Malacca, Kolkata or Jerusalem, together with a decision rule fixed in advance stating how many durable non-sacred cities would refute the exception clause.

RETIREMENT STATE AND REINSTATEMENT CONDITION

`universal_truth #25` is hereby registered as `untestable_in_place` rather than left at T2 to be quietly reused. It is not retracted — nothing here shows it false, and it may well be true. It is recorded as a claim whose only evidence is the corpus it was written from. Reinstatement to a testable state requires items 1, 2, 3 and 5 above; item 4 raises its confidence but does not by itself restore testability,

because independence of source does not repair circularity of authorship.

8 Data availability and status

The rows live in `~/ops/asia/asia.db`, tables `city` (59 rows) and `site_logic` (16 rows), opened read-only for every query in this paper. The claim under test is `universal_truth` rows 25 and 13 in `~/ops/truth.db`. The compilation the figures came from is `~/Downloads/korea-and-asian-cities.xlsx`, read with `openpyxl`; its sheets `cities`, `city_site_logic`, `summary` and `read me` are the ones cited. The absent admission rules are absent from `~/ops/ASIA-LANE-PLAN.md` §2. The loader whose docstring anticipates three of these counts is `~/ops/asia/load_cities.py`.

Status: refused. Figures pending source verification. The two confidence levels in this paper must not be collapsed. The corpus *diagnostics* in Table 1 are T1 — measured once, on 2026-09-09, from the named artefacts by the bases stated, and recomputable today by anyone holding those files. The corpus *contents* those diagnostics describe are T0 throughout: every founding year, every country and every category assignment is an unsourced assertion recorded from recall, and nothing in this paper should be read as evidence about any actual city. The medians in Table 2 are published only as the result that was declined.

Before this paper is cited: the counts should be re-derived by a reader who has not seen `load_cities.py`, and the workbook's 59-versus-60 disagreement should be resolved at source rather than banded. Before anything in the *corpus* is cited, it needs a source, which it does not currently have.

War Atlas method series, queue item 5 · dylanvir.com/war/papers · Registered 2026-09-09 · Class 3 · Published refusal.

This paper makes no claim about cities. It makes class 1 claims about three named files and one absent document. The claim registered in §2 is refused for the reasons in §5 and is not softened, reworded or partially reported; `universal_truth` #25 moves to `untestable_in_place` with the reinstatement condition stated in §7.