

Ordnance tonnage and urban defensibility: a refusal, with its search denominator

A registered study that could not be run, published as a documented absence rather than withdrawn

Measure class 2 Registered 2026-09-09 Status: **refused** — variables absent Figures pending source verification Falsifier: §7

THIS PAPER REPORTS A REFUSAL

The study described below was registered and then could not be run. No result about bombing and urban defence is offered here, and none should be quoted from this page. What is published is the absence itself — what was looked for, where, how much of the corpus was walked, and what was found in its place. Section 7 states exactly what would unblock the study.

Abstract. This study was registered to compare, across four named cities, the weight of bombs and shells delivered onto each against how much harder each then proved to take. It is refused. Neither variable exists. A sweep of 189 database tables, 74 workbook sheets and 170 battle records — 78,065 text units in all, with positive controls passing 3 of 3 — returned no figure for ordnance delivered onto any of the four places; every string resembling one resolves to a different quantity, usually supply carried *into* a place rather than munitions dropped *onto* it. The outcome side is worse: *defensibility* has no operational definition, no column and no scale anywhere in the corpus, and two of the four cases have no database record at all. Of the 24 case-by-quantity cells the study needed, 2 hold a usable value. The refusal is published rather than shelved because the claim is already circulating inside this corpus at a middling confidence tier on three rows that look like three witnesses and are one. This is a class 2 registration, so it would have made no claim about cities in general even had it run; it makes none now.

1 Background

In late August 1942 the German air force bombed Stalingrad heavily, and the city was afterwards very hard to take. A claim built on that sequence is one of the most repeated in all military writing: that the bombing did not soften the city but fortified it, because a field of broken masonry gives a defender better cover, better concealment and worse tank going than the intact buildings did. The same story is told about Monte Cassino, where an abbey was destroyed from the air and the ruins were then held; and, more quietly, about Caen and Ortona.

The claim is worth measuring because it is strategically load-bearing and almost never stated in a form anyone can check. It proposes that an input can pass a threshold beyond which more of it produces less — which, if true, is a rule about attacking cities and not an anecdote about one of them.

Before any data appears, the strongest argument against it. Rubble is not a defender's asset alone. It also breaks up the attacker's problem into small independent fights, and it did the same to the defender's command: a force fighting in a ruin loses the street plan, the telephone line and the ability to see its own flanks, exactly as its opponent does. This corpus has already ruled on that

point against itself. Its register of rejected general claims records, at entry 3, that "*terrain always favours the defender*" failed as too coarse, on the express ground that "*the same rubble that protected Chuikov's men also fragmented German command.*" The rubble is therefore already on the record here as cutting both ways, and nothing in this corpus can separate the two directions.

This paper does not settle that argument. It reports that the argument cannot be settled from anything on this machine, and states precisely what is missing so that someone else can go and get it.

2 Claim, class and preregistration

REGISTERED CLAIM

Among the cases admitted by the rule in §3, a city subjected to heavier ordnance before or during an assault proved harder to take than one subjected to less.

MEASURE CLASS

Class 2. The *measure class* records how far a finding is entitled to reach: class 1 describes a corpus, class 2 compares named cases and reaches no further, class 3 claims something about a population and needs a written rule for admitting members to it. This registration is determinate at class 2 — four named cases, no claim about cities in general — and the class was never the obstacle. A class 3 successor is not merely unwritten; it is unreachable, because a population claim would need the same two variables this paper reports as absent.

DECISION RULE, STATED BEFORE COUNTING

The rule has a gate and a test, and the gate runs first.

Gate. The study proceeds only if, for at least three of the four admitted cases, both an ordnance quantity and an outcome quantity exist with a stated *basis* — a written statement of what the number counts and how it was arrived at. If fewer than three cases carry both, the study is refused and the absence is published instead.

Test (reached only if the gate opens). The claim stands as description if the case with the heaviest ordnance shows the highest defensibility value and the case with the lightest shows the lowest, with no single case reversing the ordering. It fails if the ordering is unrelated to ordnance, or inverted.

SCALE TYPE OF EVERY VARIABLE

Ordnance would be *ratio* scale — a quantity with a true zero, on which arithmetic is meaningful. Defensibility, under any of the four operational definitions considered in §3, would also be ratio. City identity is *nominal*: a label, on which no arithmetic is permitted. The one graded column that does exist, *cover_value* in the building typology, is *ordinal* — an ordered set of words, *poor* through *excellent*, whose steps are not equal and cannot be averaged, summed or correlated. No arithmetic is performed on it anywhere in this paper, and §6 records the risk that a later worker will forget that.

A defect in the ordering must be declared here rather than buried. The corpus sweep and the class determination were carried out in a grounding pass *before* the decision rule above was written down. The rule is therefore not preregistered in the strict sense, and it is treated as a threat to validity in §6.2 rather than glossed. Its outcome is not sensitive to the ordering, because the gate turns on whether a variable exists at all — a condition with no threshold anyone could have tuned after seeing the data.

3 Admission rule, and how the set was actually found

The rule was written to be applied, and it is applied below. A case is admitted if it satisfies all six criteria.

1. The action is a deliberate attempt to take a built-up area held by a defending force, in the Second World War.
2. The place has an entity in the atlas database — a row with a name, a date and a coordinate — so that the case is addressable rather than merely nameable.
3. A quantity of ordnance delivered onto that place exists, in any unit that measures munitions arriving on the target: weight of bombs, number of sorties flown against it, or shells fired into it.
4. That ordnance quantity carries a basis distinguishing it from the several other quantities the same words are used for — supply carried in, shipping displaced, or shipping sunk.
5. An outcome quantity exists on a stated scale: days from the opening of the assault to the fall of the place, mean advance in metres per day, or the attacking force's share of loss.
6. A comparison point exists — either the same ground measured before the bombardment, or a matched place of similar built form that was not bombarded — so that *harder to take* is measured against something.

Applied. Criterion 1 admits all four candidates. Criterion 2 admits two: Stalingrad and Monte Cassino have rows in the atlas battle table; Caen and Ortona have none. Criterion 3 admits none. Criteria 4, 5 and 6 are therefore never reached, and no case is admitted. The gate in §2 closes at zero of a required three.

How the set was found, in plain words. Not by rule. The four cases were named in the method series' own queue, where they appear as the four famous rubble battles, and they entered this study on that basis. Two of them were named before anyone checked whether they exist as data. That is a selection effect of the plainest kind and it is recorded again in §6.1.

4 Data

Table 1 is the complete inventory of what the sweep returned. It is not a table of the study's variables, because the study has none; it is a table of everything on this machine that a reader might reasonably expect to serve as one, with the reason each cannot. Every row carries two required columns.

Basis states what the quantity counts and how it was arrived at, because the same word covers different quantities: tonnage carried in is not tonnage dropped on, and a rate per gun per day is not a total. **Confidence** takes one of four values, defined here and used consistently:

- **none recordable** — the table holding the row has no field in which a source or a confidence could be written, so the row cannot carry provenance even in principle.
- **T0 scaffold** — present in a reference workbook whose own read-me states that the file was assembled from recall without a citation trail. *T0* is the lowest tier in this estate's grading: an assertion whose only support is that someone wrote it down.
- **unverified** — the row names a source reference and its own status field reads `unverified`. Better than *T0* in that it is addressable; not yet checked against anything.
- **raw, unparsed** — a verbatim string is stored, but the numeric fields beside it are empty and the row is flagged as not parsed, so nothing arithmetic can be done with it.

Table 1. Everything the sweep returned that bears on the question, with the reason none of it is the variable. Fourteen items. Not one is a quantity of ordnance delivered onto a city. Rows 13 and 14 were found during re-verification for this paper and correct the grounding inventory, which had reported no sortie count and no ordnance-expenditure rate anywhere.

ITEM	WHAT IT SAYS	QUANTITY AND UNIT	BASIS	CONFIDENCE	ADDRESS
------	--------------	----------------------	-------	------------	---------

ITEM	WHAT IT SAYS	QUANTITY AND UNIT	BASIS	CONFIDENCE	ADDRESS
1 · Rubble field	"the paradox: bombing improves the defence"; behaves as "cannot be destroyed further"	none; cover_value "excellent" (ordinal)	A cell in a 12-row typology of building kinds. Ordinal grade over the typology, not a measurement of any city. No join object connects the typology to any battle or place. Typical setting given as "Stalingrad after August 1942".	T0 SCAFFOLD	material-layer.xlsx [building_types] r12
2 · Cellar	"below the depression limit of a tank gun"; "Stalingrad's decisive layer"	none; cover_value "excellent" (ordinal)	Same typology, same absent join. A statement about a building kind, not a count of cellars in any city.	T0 SCAFFOLD	material-layer.xlsx [building_types] r9
3 · Grain elevator	"near-immune to field artillery"; "height plus concrete plus interior galleries"	none; cover_value "excellent" (ordinal)	Same typology. Names one structure type at Stalingrad; carries no count, no ordnance figure and no outcome.	T0 SCAFFOLD	material-layer.xlsx [building_types] r11
4 · Reinforced concrete	"the best urban cover; rubble stays in blocks"	none; cover_value "excellent" (ordinal)	Same typology. Five of its twelve rows are graded plainly "excellent" and two more carry qualified variants of it, so the grade barely discriminates within the typology, let alone between cities.	T0 SCAFFOLD	material-layer.xlsx [building_types] r6
5 · Mechanism M-RUBBLE	"Rubble defends better than intact buildings"; note: "the bombing meant to break the city built the terrain that held it"	none	Unattributed assertion. The table holding it has five columns — id, mechanism, inverts, why, note — and neither a source field nor a confidence field, so no provenance can be attached to it without altering the schema.	NONE RECORDABLE	atlas.db stal_mechanism M-RUBBLE
6 · Fragment density	Metal fragments recovered per unit area at one hill	1,250 fragments per m ²	Density recovered from Mamayev Kurgan soil, dated post-war. A measurement of the ground, not of the battle — the row's own note says so. Nearest thing on this machine to an ordnance-intensity figure, and it has no comparator at any other place.	UNVERIFIED	atlas.db stal_figure F-FRAG
7 · Airlift thresholds	Four daily tonnages: bare survival, survival plus defence, full combat, and what was actually delivered	300 · 500 · 750 · 100 tons/day	Wrong quantity and wrong direction. Supply flown <i>into</i> the encircled pocket, not munitions delivered onto a city. Treating these as ordnance would be exactly the basis error the estate's own register warns of.	NONE RECORDABLE	atlas.db stal_airlift , 4 rows
8 · Monte Cassino entity	Battle record: 17 Jan – 18 May 1944, 41.490 N, 13.8139 E	dates and coordinate only	Admitted from an external identifier under rule ADMIT-WD-002. Carries zero side rows, and therefore no strength, no casualties and no outcome of any kind.	UNVERIFIED	atlas.db battle Q3 35288

ITEM	WHAT IT SAYS	QUANTITY AND UNIT	BASIS	CONFIDENCE	ADDRESS
9 · Stalingrad entity, twice	Two battle rows for one battle, under two different admission rules	dates 1942-08-23 →1943-02-02 on one row	Q38789 carries the dates and coordinate but no sides; the legacy row carries the only two side rows but no dates or coordinate. A count over the battle table double-counts the case, so the case count itself is unstable before any measurement begins.	UNVERIFIED	atlas.db battle Q38789 · battle-of-stalingrad
10 · Stalingrad strengths	"~1,000,000+ (Stalingrad, Don, and Southwestern Fronts)" and "~1,000,000 (German 6th Army and Axis allied armies)"	no numeric value; low, high, unit all null	Verbatim source strings. Flagged unparsed, scale structural, measure moment unknown — so it is not even recorded whether these are strengths at the start, the peak or the end. Two figures from one compilation are in any case one witness, not two.	RAW, UNPARSED	atlas.db strength, sides 81 & 82
11 · Pocket strength	Axis personnel encircled at the closing of the ring	290,000 personnel as of 1942-11-23	A population trapped, not a defence outcome. The row's own note warns that "trapped" and "German" are different populations, the pocket having held Romanian and Croatian formations and Soviet auxiliaries.	UNVERIFIED	atlas.db stal_figure F-POCKET
12 · Stalingrad casualties	"Hundreds of thousands killed/wounded; ~91,000 captured"; "Over 1,000,000 total casualties"	prose; confidence field reads "disputed"	Casualties are not deaths, and these strings mix killed, wounded, missing and captured against unstated denominators. The record's own note reads "casualty estimates vary widely across sources". Monte Cassino is absent from this file entirely.	RAW, UNPARSED	battles.json, battle-of-stalingrad
13 · Ju 52 sortie count	Sorties required to deliver the survival threshold	172 sorties	The only sortie count on this machine, and it is not ordnance. A spreadsheet formula, <code>=ROUNDUP(300/1.75,0)</code> , dividing the workbook's own 300 t/day threshold by the midpoint of its own 1.5–2 t per-sortie band. Transport flying supply in, computed from two other cells of the same file. It is arithmetic on a scaffold, not an observation.	T0 SCAFFOLD	command-and-tactics.xlsx [summary] r27
14 · Artillery expenditure rate	Rounds fired per gun on an active day	20–100 rounds per gun per active day	The only ordnance-expenditure figure on this machine. A doctrine-generic band with no gun count, no count of active days and no attachment to any place — the note concedes it "varies enormously by phase and doctrine". It cannot be multiplied into a total for any of the four cases without inventing both other factors.	T0 SCAFFOLD	command-and-tactics.xlsx [logistics_rates] r18

Figure 1. *The shape of the absence.* Six quantities the study required, across the four named cases: 24 cells, of which 2 hold a usable value and 21 are empty. Every cell in the ordnance row, the outcome row, the baseline row and the built-form row is empty for every case. **No trend line, curve or fitted relationship is drawn, and none could be:** a line requires two variables with values, and neither variable here has a single value at any case. The reason is stronger than the usual class 2 reason — it is not that four points are too few to fit, but that there are no points. Counts recomputable by the command in §8.

5 Results

The decision rule's outcome: refused at the gate. The gate in §2 required both variables present, with a stated basis, for at least three of four cases. The count is zero. The test was never reached and no ordering was examined, because there is nothing to order.

The *search denominator* — the size of what was actually looked through, without which a negative result is only a claim about someone's diligence — is 189 sqlite tables across three databases, 74 sheets across eight reference workbooks, and 170 battle records, comprising 78,065 text units. The sweep carried three positive controls, terms known in advance to be present, so that a sweep matching nothing could be distinguished from a sweep that was broken: *rubble* returned 14 hits, *Stalingrad* 59, *Monte Cassino* 2. All three passed. *Ortona* returned zero across the entire corpus. *Caen* returned one, a prose clause about the Normandy campaign inside a single record about the D-Day landings.

Eight strings did match the ordnance pattern, and each was read by hand rather than counted. All eight are the same word doing different work: *sortie* in its older sense of a sally out of a besieged place, at Carrhae, at the siege of Rome and at Panipat; a warship sortieing from harbour before Midway; *tonnage* meaning shipping sunk at the Coral Sea, merchant shipping committed to Normandy, and a rate of carrier construction. Two further items, rows 13 and 14 of Table 1, are the closest true approaches and both fail on basis: a count of transport sorties flying supply inward, derived by formula from two other cells of the same workbook, and a generic band of rounds per gun per active day attached to no place, no gun count and no number of days.

The honest reading. Nothing in this corpus supports the claim, and nothing in it refutes the claim. Those are different from each other and both are different from the claim being false, and the distinction is the whole content of this paper. What can be stated positively is narrower and more uncomfortable: the claim is already circulating *inside* this corpus as though it had been established, and it has not been.

The mechanism row asserts that the bombing meant to break the city built the terrain that held it, and sits in a table with no field for a source. The building typology states, at the row for a rubble field, that bombing improves the defence. The register of general claims elevates the same sentence to tier T2 — a claim held to have passed an internal consistency check — and cites Stalingrad as its lead instance, while its own sourcing field reads: "*recall plus the 2026-09-09 conversation. No citation trail, no fresh research... not publishable without a named source.*"

Three tables, three files, three apparently converging lines. All three carry the same date, all three were written by the same hand in the same pass, and two of them declare in their own metadata that their basis is recall. Under this estate's rule that independence is *extractional* — that agreement counts as corroboration only when the agreeing statements travelled by genuinely different routes — that is one attestation in three coats, not three attestations. Had this paper been written to its brief, it would have footnoted those three rows beside one another, and would thereby have laundered a single unsourced belief into a corroborated finding. That is precisely the failure the atlas exists to catch, committed by the atlas, and it is the reason the refusal is worth publishing rather than shelving.

The row that fits worst is Monte Cassino, and it fits worst in a way that ought to be uncomfortable. It is the one case other than Stalingrad with a database entity, so it is the case that would have carried the comparison — and that entity is empty. It has a name, two dates and a coordinate, zero side records, and therefore no strength, no casualties and no outcome of any kind. It is also the case where the popular version of the claim is at its most vivid: an abbey bombed to rubble and the rubble then held. The case whose story does the most work for the thesis contributes the least evidence to it. With Cassino empty and Caen and Ortona absent, the study reduces to Stalingrad, and one case is not a comparison; it is the anecdote the study was registered to test.

One further figure deserves naming precisely because it is the sort of thing that gets pressed into service. The corpus records that one building at Stalingrad, Pavlov's House, was held for 58 days. It is the only duration-of-defence figure on this machine, it carries a source reference and its own status field reads *unverified*, and it measures one building rather than a city. It is not an outcome variable for a case and must not be used as one.

6 Threats to validity

These threaten the refusal, which is the only claim this paper makes.

- 1 **Selection.** The four cases were chosen because they are the famous rubble battles, and famous partly because the story attaches to them. Cities bombed heavily and then taken quickly are less written about and would never have been found by the method used here. Had the gate opened, this alone would have capped the result at description.
- 2 **The rule was written after the search.** The corpus sweep and the class determination preceded the written decision rule in §2. The rule is therefore not preregistered in the strict sense. Its gate turns on presence rather than on a threshold, which limits but does not eliminate the concern.
- 3 **Absence in this corpus is not absence in the world.** Bomb tonnages for these operations plainly exist in published sources. This paper reports that none of them has been brought onto this machine, which is a statement about the corpus and about nothing else. Reading it as a claim that the figures are unobtainable would be a serious misuse.
- 4 **One reader, one pass, one method.** The sweep, the hand-reading of its eight matches and this report were all produced by a single agent in a single session. By the estate's own independence rule that is one extraction path, not a replication, and the positive controls — which passed — were written by the same hand that wrote the search.
- 5 **The negative may be a naming problem.** The sweep matches strings. A tonnage stored as a bare number in a column named something unanticipated would be missed by the term list, though the walk over every column name in all 189 tables reduces that risk. It does not eliminate it, and criterion 3 of §7 is written to catch exactly this.
- 6 **Every catalogued failure is one that was caught.** The eight false matches were found and read. Whatever the pattern did not match was neither found nor read, so this account understates the corpus's silence-and-noise by an unknown amount in an unknown direction.
- 7 **The ordinal column is an invitation.** *cover_value* is a word grade over twelve building kinds, five of which read plainly "excellent" and two more a qualified version of it. It is the only graded column in the vicinity, and the standing risk is that a later worker converts those words to numbers, correlates them with something and reports a result — reversing this refusal without adding a single measurement. If that appears, it is not a successor to this paper; it is the failure this paper was written to prevent.
- 8 **The case count itself is unstable.** Stalingrad occupies two battle rows under two admission rules, which disagree about which of them carries the sides. Any count over the battle table double-counts the case before measurement begins.

7 What would refute this

The claim being retired here is not *bombing fortified the city* — that claim was never established and so cannot be retired. What is registered is the refusal, and the refusal is checkable by someone who has none of this data.

- A named published source giving a quantity of ordnance delivered onto any two of the four places, with a stated basis and a stated period, in any unit that measures munitions arriving on the target. Two cases reopen the comparison; three reopen the gate.
- A published operational definition of urban defensibility, with values for at least two of the four on the same definition — days from the opening of an assault to the fall of the place, mean advance in metres per day, or attacker loss share, each against a stated denominator.

- A second, extractionally independent reading of this same corpus, by a different person using a different method, that returns an ordnance quantity the sweep in §8 missed. This refutes the absence without touching the underlying claim, and it is the most likely of the three to succeed.
- A pre-bombardment and post-bombardment pair on the same ground, or a matched comparator city of similar built form that was not bombarded, on either of which the outcome variable is measured. Without one of these the study has a treated case and no control, and would fail its own positive-control requirement even with both variables in hand.
- A join object connecting the building typology to any city — a table giving the share of a named place by construction kind. Without a denominator of built form, no one can say a city became more defensible; only that a rubble field is good cover in the abstract.

Retirement state. The registered claim moves to `refused_variable_absent`, distinct from *not supported*: nothing was tested, so nothing failed. **Reinstatement condition:** both variables present with a stated basis for at least three of the four cases, under the §3 rule unchanged, plus one comparison point satisfying criterion 6.

Three consequences follow inside the corpus and are recorded here as owed work rather than performed. First, the mechanism table needs a source field and a confidence field before any row in it is cited again; at present it cannot carry provenance even if someone wanted to add it. Second, the tier-T2 general claim that cites Stalingrad as its lead instance is not adjudicated by this paper — its broader statement about inputs may well survive — but that particular instance is not attested here and should not be quoted as though it were. Third, none of the nine attested-absence records in the atlas covers this gap; all nine concern ancient battles — commanders, identities and quantities at Ur, Kadesh, Lachish, Changping, the Hydaspes and Red Cliffs — so the machinery for recording a documented absence exists and has never once been pointed at a modern case. One should be opened against this absence, carrying the search denominator and the positive-control result from §5, so that the next person to ask this question finds the answer already recorded rather than repeating the sweep.

8 Data availability and status

Every item in Table 1 is addressed in its final column and can be read directly. The databases are `war-atlas/db/atlas.db`, `ops/truth.db` and `ops/asia/asia.db`; the workbooks are the five in `ops/reference/` plus three further research workbooks; the battle records are `war-atlas/data/battles.json`. All were opened read-only.

The search denominator, the positive controls and the 24-cell coverage count in Figure 1 are not described here, they are recomputed by one command:

```
python ~/ops/papers/ordnance-sweep.py
```

It walks every table in the three databases and every sheet in the eight workbooks, prints the denominator, runs the three positive controls with their expected minimum hit counts asserted, lists every match for hand-reading, and derives the coverage matrix from the live database rather than from anything written down. As of 2026-09-09 it reports 189 tables, 74 sheets, 170 entries, 78,065 text units, positive controls 3 of 3, and coverage of 2 usable cells, 1 raw-text cell and 21 empty cells of 24. A run that disagrees with any of those figures supersedes this paper's version of them.

Status: refused; figures pending source verification. The two status lines are separate and both apply. *Refused* is the paper's result and is not provisional. *Figures pending source verification* applies to every quantity in Table 1 without exception: four are T0 scaffold from workbooks whose own read-me states they were assembled from recall, four carry a source reference with status `unverified`, two are raw unparsed strings, and two sit in a table with nowhere to record a source at all. Not one figure in this paper has been checked against a named external source, and none should be quoted from here. Publishing at this confidence is deliberate: the contribution is the documented absence and its denominator, and both are improved by someone bringing a real tonnage rather than by waiting.

War Atlas method series, paper 02 · dylanvir.com/war · Registered 2026-09-09 · Class 2 · **Refused for want of both variables.** This paper makes no claim about bombing and urban defence, and no population claim about cities. A class 3 successor is not merely unwritten but unreachable: it would need the same two variables reported absent here. Recompute: `python ops/papers/ordnance-sweep.py`.