Epistemic Enclosure

A COMPUTATIONAL-STATUTORY AUDIT OF ENGLISH PLANNING

RESEARCH QUESTION

What are the structural and informational conditions under which planning contestation occurs in England, and how are they distributed across socioeconomic gradients?

This thesis investigates the structural and evidentiary asymmetries within the English statutory planning system. It deploys a federated spatial data pipeline that programmatically cross-references contested planning sites against 65 live open-government spatial layers, complementing existing qualitative case-study research in this field. The empirical foundation is a locked, audited cohort of N = 652 contested planning sites, drawn from a gross harvest of N = 852 grassroots campaigns via a seven-stage sample-flow attrition funnel followed by collapsing repeat submissions of the same site.

Empirical foundation	Audited, de-duplicated cohort N = 652 (strict v6; from N = 852 gross harvest)
Computational infrastructure	Federated spatial pipeline; 33 APIs; 65+ statutory layers; 7-stage attrition funnel + repeat-submission de-duplication
Finding 1 — Defensibility Gap (two axes)	Ancient Woodland skews affluent (OR = 4.88, p = 8.4×10^{-6}); composite score with flood removed U = 8,370.5, **p = 0.0017**; Heritage at Risk skews deprived (OR = 7.54, p = 3.2×10^{-5})
Finding 2 — BMV farmland	$\chi^2 = 12.72$, df = 3, p = 0.0053 (n = 652); energy 72.7%, sprawl 54.2%, logistics 50.0%, infill 46.6%
Finding 3 — Ecological data scarcity	34.9% data deserts; no detectable class difference (p = 0.34)
Finding 4 — State API fragility (three services, three failure modes)	Probes of 4 Oct 2026, three runs, one vantage point. HM Land Registry INSPIRE: layered refusal — scripted requests refused at the Cloudflare edge (403); a manual browser capture returned `Service WFS is disabled`. HSE MAHP zones: no national service found by the pipeline; two council layers queried (absence is not testable by probe). Coal/MRA: configured addresses invalid (HTTP 200 carrying an "Invalid URL" error); the layer is discoverable through the public catalogue and answers

Theoretical claim: Epistemic Enclosure — domain extension of Tseng, Becker & Roikonen (2024), not a coinage
Contribution: Complementary computational-statutory audit of the evidentiary preconditions of contestation

Phase-2 cohort extension

- revealed car dependency 76.6% (TS061, correct LSOAs, all 652 sites) ·
- Heritage at Risk 5th shield (OR = 7.54, p = 3.2×10^{-5}) ·
- Brownfield register (only 4.8% registered; 93.1% of infill not on the register) ·
- Land ownership: 39.1% unresolved corporate ownership via postcode approximation (Mechanism A; v5: 309 classified corporate entities, 7 STRICT housebuilders) ·
- Outcomes: 90.1% permitted (desert→permit OR = 0.98, clean null) ·

- **Track E:** Sprawl under-represented in data deserts (28.5% vs 42.1%, OR = 0.55, p = 0.00035) ·
- **Track F**: 12,824 GM empty homes equal 61.0% of PfE Green Belt target (21,030 units) ·
- **Track G:** 4.5× borrowing wall breached in all 10 GM boroughs (5.39×–8.81×)

SAMPLE PROVENANCE — 7-STAGE FUNNEL + DE-DUPLICATION

The pipeline enforces a deterministic, auditable sample-flow funnel. No metric in the thesis may cite an unverified denominator.

The funnel

Stage	*Description*	*N Before*	*Excluded*	*N After*
1	Gross harvest seed cohort	—	—	852
2	English statutory jurisdiction filter	852	−35	817
3	IMD Health Deprivation Domain record present	817	-25	792
4	**LSOA resolved (geocoded)**	**792**	**-49**	**743**
5	**LSOA matched to 2019 IMD File 7 boundaries**	**743**	**-42**	**701**
6	**Pipeline A (Site Intelligence Summary) present**	**701**	**0**	**701**
7+8	**Pipeline C present; repeat submissions collapsed**	**701**	**-49**	**652**
—	**Final canonical cohort**	**—**	**—**	**652**

Pipeline A and Pipeline C each matched all 701 pre-deduplication records (100%). The strict v6 cohort is N = 652 after collapsing 49 repeat submissions.

Canonical stratification (N = 652)

Stratum	*N*	*Share*
High Deprivation (IoD2025 1–3)	87	13.3%
Intermediate Deprivation (IoD2025 deciles 4–7)	321	49.2%
Low Deprivation / Affluent (IoD2025 deciles 8–10)	242	37.1%
Unresolved IMD decile (boundary shift at LSOA level)	2	0.3%
Triple Jeopardy (canonical cohort)	47	7.2%
Triple Jeopardy (gross harvest)	54	6.3%

On sample bias: high-deprivation sites (IMD 1–3) represent the more capable fraction of disadvantaged wards — a convenience sample structurally under-represents communities without organising capacity. The likely direction of selection bias is conservative: the true national gap is plausibly larger than what is measured here, not smaller.

TEN CORE STUDIES, THREE ROBUSTNESS CHECKS

#	*Study*	*Key statistic*	*Conclusion*	*Level*	*N*
1	**Defensibility Gap — ecological axis**	U = 8,370.5; p = 0.0017 (flood removed); Ancient Woodland OR = 4.88 [2.25, 10.58], p = 8.4×10^{-6}	Statutory protection scales with wealth — on the ecological axis	L2	652 (329 in extremes)
2	**Defensibility Gap — heritage counter-axis**	Heritage at Risk OR = 7.54 [2.80, 20.34], Fisher p = 3.2×10^{-5}; five-shield composite v5 p = 0.14 (N = 641), not yet recomputed on N = 652	Protection runs on two axes; the gap is in affluence-linked preventative shields, not protection per se	L2	652 (LOW_N: affluent cell n = 6; Fisher p = 3.2×10^{-5})
3	**BMV farmland × development archetype**	$\chi^2(3) = 12.72$, p = 0.0053 (n = 652); energy 72.7%, sprawl 54.2%, logistics 50.0%, infill 46.6%	Speculative infrastructure concentrates on high-grade agricultural land	L2	652
4	**Ecological Data Deserts**	34.9% of sites in data deserts; p = 0.34 (null class result); IoD2025 decile 6 peak at 47.4%	Data scarcity is an infrastructural condition of the state — non-monotonic, peaking in the peri-urban fringe; not class-targeted	L1 / L2 null	639
5	**State API fragility**	HM Land Registry INSPIRE: scripted requests returned 403 at Cloudflare; manual browser returned “Service WFS is disabled”. HSE: no national service found by the pipeline. Coal/MRA: configured URLs returned HTTP 200 with “Invalid URL” errors, while a catalogue-discovered layer responded.	Key statutory data is open in law, closed in operation — the practical gap is where enclosure bites	L1	Direct observation
6	**Car dependency**	76.6% car/van commute share (mean); 35.0% WFH	Peripheralisation is behavioural, not merely infrastructural — access ≠ substitution	L1	652 sites (Census 2021 TS061; correct LSOAs)
7	**Brownfield "infill" label**	31/652 on register (4.8%); infill archetype 6.9% vs others 3.6%; $\chi^2(3) = 5.52$, p = 0.14	The "infill" label is not supported by registration — 93.1% of infill-labelled sites are not on the register	L1 / L2 null	652
8	**Land opacity (Christophers, Mechanism A)**	255/652 (39.1%) unresolved corporate ownership via postcode approximation	The registration gap is measurable inside the cohort — converts a citation into a measurement	L1	652 (397 matched corporate-owned title; 255 unresolved via postcode approximation)
9	**Outcome prediction (PlanIt + PINS)**	606 decided sites with resolved NBN count: 546 permitted, 60 refused (90.1%); desert↛permit OR = 0.98 [0.56, 1.71], Fisher p = 1.00; unresolved ownership OR = 0.75 [0.44, 1.28], p = 0.34 (n = 619 decided)	The consultation-stage asymmetry does not translate into a detectable difference in the binary permit outcome — a clean null that bounds the thesis	L1 for records / L2 tested null	650/652 sites resolved; 205/206 LPAs (PINS)
10	**Triple Jeopardy × Labour seats**	41/47 (87.2%) of canonical Triple Jeopardy sites Labour-held vs 41.8% non-TJ; OR = 9.51 [3.98, 22.73], Fisher p = 7.4×10^{-10}	Compound vulnerability concentrates in Labour-held constituencies, but with a clean baseline and a tested association	L2	47 TJ / 605 non-TJ

THE EVIDENTIARY PROBLEM IN ENGLISH PLANNING

Extended Urbanisation ↠ 2. Epistemic Enclosure ↠ 3. Inscriptive Injustice ↠ 4. Asymmetric Statutory Enclosure

Concept	*Source*	*Empirical manifestation*	*Evidentiary status*
Statutory planning procedure	TCPA 1990 / NPPF	21-day consultation window	Level 1 (direct observation)
Extended urbanisation	Brenner & Schmid, 2014; Schmid & Topalović, 2023	Peri-urban operational landscapes	Level 3 (plausible mechanism)
Inscriptive injustice	Parker et al., 2026	Community evidence discounted as anecdotal	Level 2 (statistical association)
Epistemic Enclosure	*this study*, 2026 — domain extension of Tseng et al., 2024	Evidentiary asymmetry during consultation	Level 3 (plausible mechanism)

Statutory planning determinations in England are fundamentally evidentiary procedures governed by the Town and Country Planning Act 1990 and the National Planning Policy Framework. Commercial applicants mobilise multidisciplinary consultancies to produce Environmental Impact Assessments, transport dispersion models, hydrological flood risk assessments, and Biodiversity Net Gain metrics. Host communities, absorbing the localised externalities of these developments, operate under a statutory consultation window of just 21 days with no public funding for counter-expertise.

CONCLUSION

The English planning system distributes its evidence, its protections, and its data unevenly — and the unevenness tracks wealth for the kinds of protection that matter most, while leaving everyone equally exposed at the moment developers apply.

#	*Headline Conclusion*	*What It Means in Plain English*	*Key Empirical Anchor*	*Evidence Level*
1	The system protects wealth, not need	Affluent areas are heavily shielded by statutory designations that legally halt development, while deprived areas receive far fewer shields and bear more hazard burdens. This inherited geography directly determines present-day legal weight when developers apply.	Ancient Woodland odds 4.88× higher in affluent areas ($p = 8.4 \times 10^{-6}$); flood-removed composite U = 8,370.5, p = 0.0017.	Level 2
2	Protection runs on two opposite tracks	Protection is not a single wealth axis. Nature and tree shields cluster in wealthy areas, while historic built-heritage designations cluster in poorer, post-industrial areas. The system protects preservation where wealth is, and decay where deprivation is.	Heritage at Risk odds 7.54× higher in deprived areas ($p = 3.2 \times 10^{-5}$); five-shield composite v5 p = 0.14 (N = 641), not yet recomputed on N = 652.	Level 2
3	Data deserts belong to infrastructure, not sprawl	Housing sprawl does not hide in unmapped biological voids—utility-scale solar, battery storage, and freight logistics hubs do. For housing sprawl, enclosure happens through secret land titles, prime soil conversion, and the 21-day clock.	Sprawl is under-represented in data deserts (28.5% vs 42.1%, OR = 0.55, p = 0.00035); energy (63.6%) and logistics (47.1%) have the highest desert rates.	Level 2
4	Developers target the best farmland	Speculative infrastructure and peripheral expansion heavily target high-grade agricultural land rather than reusing brownfield plots within existing urban footprints.	72.7% of energy, 54.2% of sprawl, 50.0% of logistics and 46.6% of infill sites intersect provisional BMV soils ($\chi^2 = 12.72$, p = 0.0053).	Level 2
5	The state actively encloses its own data	Government data infrastructure is not just "broken"; it is withheld by design under three distinct institutional motives: national security, commercial copyright disputes, and state search-fee monopolies.	HSE: no national service found by the pipeline; HMLR: scripted requests returned 403; manual browser returned “Service WFS is disabled”; Coal/MRA: configured URLs returned “Invalid URL” errors, while a catalogue-discovered layer responded.	Level 1
6	Contesting early does not change the binary permit outcome	Facing a data desert or opaque landowner does not change whether a developer gets permission. The deck is procedurally stacked against communities inside the 21-day window, but developers win roughly 90% of applications regardless.	PlanIt outcomes: 606 decided sites; 546 permitted (90.1%); data-desert predictor null (OR = 0.98, p = 1.00).	Level 1 (permits) / Level 2 (null)
7	The "land shortage" is an illusion	Councils justify bulldozing the Green Belt on the premise that housing land is exhausted. In reality, existing long-term empty dwellings alone could absorb the majority of statutory 15-year targets without touching a single field.	12,82412,824 long-term empty homes = 61.0%61.0% of the entire 15-year Greater Manchester Green Belt plan (21,03021,030 units); Bolton (157.8%157.8%) and Salford (207.8%207.8%) exceed plan.	Level 1
8	New peripheral builds do not help the young	De-risking greenfields for £400k+ executive commuter homes does not trickle down to the next generation. Young families are locked out by mortgage borrowing caps and trapped renting cold, 120-year-old Victorian terraces that receive zero state investment.	Lower-quartile price-to-earnings breaches the 4.5×4.5× cap conurbation-wide (5.39×5.39× to 8.81×8.81×); working-age private renting is 37.6%37.6% in Manchester, while pre-1919 stock is 30%–35%30%–35%.	Level 1 (ratios) / Level 3 (mechanism)
9	Land ownership is three times more secretive where contested	While almost 90% of England and Wales is registered on official land records, cadastral opacity spikes dramatically at contested sites. Communities are forced to contest schemes against invisible corporate proprietors and shell companies.	39.1% unresolved corporate ownership via postcode approximation (255/652 sites); this is not a measure of unregistered land.	Level 1

These conclusions are grounded in a rigorous computational audit: a federated spatial pipeline that cross-references 652 contested planning sites against 65 live government datasets, with a formal evidence-ladder discipline that separates direct observation from statistical inference and makes no causal claims.

EXTENDED URBANISATION AND EPISTEMIC ENCLOSURE

Key formulation: Epistemic enclosure is proposed as the micro-scalar mechanism operationalising extended urbanisation — a theoretical framing offered as an interpretive lens for the empirical findings, not itself an empirical claim.

Theoretic anchor	*Concept*	*Domain*	*Distinction from Epistemic Enclosure*
Lefebvre, 1991	Representations of space (conceived space)	Philosophy of space	Epistemic Enclosure is the micro-scalar mechanism of conceived space
Tseng, Becker & Roikonen, 2024	"Epistemic enclosures" (plural)	Digital participatory budgeting (Helsinki)	Domain extension: baseline evidence generation vs. proposal evaluation
Harvey, 2003	Accumulation by dispossession	Material/economic alienation	Epistemic Enclosure is the informational prerequisite
Kitchin, 2014	Data friction	Data movement and standardisation	Epistemic Enclosure is the timeline-scarcity conjunction
Mattern, 2021	Urban dashboards	Hyper-visibilisation of metro core	Epistemic Enclosure is the geographical inverse: peripheral invisibilisation
Parker et al., 2026	Inscriptive injustice	Neighbourhood Plan inscription failure	Pushes frontier to statutory API and baseline-evidence layer
Schmid & Topalović, 2023	Planetary struggles	Qualitative mapping of 8 global regions	Complementary computational-statutory audit

Extended urbanisation reframes the traditional city/countryside dichotomy by positing that urbanisation operates as an omnipresent process restructuring terrain to support metropolitan metabolisms. This chapter positions Epistemic Enclosure as a domain extension of Tseng et al. (2024) rather than a coinage.

Phase-2 — Revealed car dependency (Census 2021 TS061, correct LSOAs across all 652 sites): mean 76.6% car/van commute share; 35.0% work mainly from home (read as a ceiling given pandemic-era enumeration). Access is not the same as dependency — these operational landscapes show near-universal car reliance among those who travel, grounding the peripheralisation claim in behaviour, not just infrastructure. Level 1, direct observation.

THEORETICAL FRAMEWORK: THE LINEAGE OF SPATIAL JUSTICE

This chapter establishes the thesis's theoretical scaffolding. It bridges critical urban theory, regenerative design, and commons governance to diagnose the structural asymmetry of the planning system.

Theoretical Tradition	*Key Theorists*	*Diagnostic Contribution to this Thesis*
The Production of Space	Henri Lefebvre (1991)	The Algorithmic Production of Abstraction: API failures and data deserts are not glitches; they are the triumph of "conceived space," rendering complex ecosystems as blank administrative slates.
The Spatial Fix	David Harvey (2003)	Speculative Enclosure: Surplus capital seeks unallocated greenfields to bypass the friction and cost of urban brownfield renewal.
Relational Place	Doreen Massey (2005)	Community as Multiplicity: Place is the "throwntogetherness of stories-so-far." Community input must be captured as structured, multi-vocal evidence, not romanticized sentiment.
Commons Governance	**Elinor Ostrom (1990)**	**Civic Infrastructure Ownership: A statutory measurement instrument built to defend communities against asymmetric power cannot reproduce that asymmetry in its own ownership.**
Inscriptive Injustice	***Parker et al. (2026)***	***The Defensibility Gap: Disadvantaged communities lack the pre-existing legal designations required to have their lived realities written into statutory policy.***
Theoretical Tradition	Key Theorists	Diagnostic Contribution to this Thesis

key formulation: "Epistemic enclosure is the micro-scalar mechanism of conceived space that operationalises extended urbanisation, translating data absence into a statutory presumption of low environmental value."

LITERATURE POSITIONING: WHAT THIS STUDY ADDS

Epistemic enclosure is the micro-scalar mechanism of conceived space that operationalises extended urbanisation, translating data absence into a statutory presumption of low environmental value.

Wave	*Period*	*Method*	*Representative scholar*	*Blind spot*
Wave 1	2014–2018	Macro-cartographic abstraction	Katsikis (2016)	Scalar abstraction: maps the "what," not the "how"
Wave 2	2018–2025	Ethnographic and political ecology	Arboleda (2020); Menon (2025); Schmid & Topalović (2023)	Lacks statistical generalisability at national scale
	2026	Computational spatial analytics + statutory auditing	***This Study (2026)***	Comprehensive void: synthesises critical urban theory with data science and statutory law

Benchmarking matrix — key comparisons:

Scholar & Work	*Theoretical linchpin*	*Empirical scope*	*Methodology*	*Blind spot addressed*
Katsikis (2016)	Operational landscapes	Planetary	Macro-cartographic visualisation	Scalar abstraction
Arboleda (2020)	Supply chain capitalism	Atacama Desert	Qualitative political ecology	Quantitative rigour
Angelo (2021)	Urbanized nature	Ruhr Valley	Historical sociology	Statutory mechanics
Schmid & Topalović (2023)	Planetary struggles	8 global regions	Qualitative mapping	Statistical generalisability
Tseng et al. (2024)	Epistemic enclosures	Helsinki	Platform UX analysis	Domain extension to statutory baseline
This Study (2026)	Epistemic Enclosure	England; N = 652	Federated spatial pipeline, χ^2, Log-Defensibility Ratio	Comprehensive void

This chapter maps the academic landscape that extended urbanisation research has generated between 2014 and 2026. The literature is categorised into three methodological waves: Wave 1 (macro-cartographic abstraction), Wave 2 (ethnographic and political ecology turn), and Wave 3 (computational spatial analytics and statutory auditing, introduced by this thesis). The benchmarking matrix positions each seminal work against the blind spot this thesis addresses.

RESEARCH DESIGN: THE FEDERATED SPATIAL PIPELINE

This chapter documents the federated software architecture executing 33 parallel asynchronous HTTP requests against 65+ live open-government spatial layers. A key technical challenge is the coordinate projection gap: statutory UK environmental constraints are natively indexed in British National Grid (EPSG:27700), whereas standard web APIs output WGS84 (EPSG:4326). Passing WGS84 coordinates directly into a statutory endpoint casts query points into the Atlantic Ocean, producing false negatives across the entire cohort. An in-memory geodetic engine was built in pure TypeScript to resolve this.

Pipeline component	*Function*	*Technical specification*
Federated orchestration	Parallel querying of 33 APIs	Bounded concurrency, max 1–4 workers, 250ms jitter
Coordinate reprojection	WGS84 ↔ British National Grid	Airy 1830 ellipsoid, Transverse Mercator, Taylor expansion to 6th power
Buffer envelopes	Statutory proximity queries	500m Ancient Woodland; 1,000m NBN / SSSI; 200m ALC soil
Data ingestion sources	**8 ministries/agencies**	**MHCLG, EA, Natural England, Historic England, DfT, Coal Authority, HSE, ONS/NHS**
Output schema	***Typed structured records***	***HazardStatus ∈ {TRUE, FALSE, ERROR, OUT_OF_BOUNDS}***

Geodetic constants (Airy 1830 ellipsoid):

Parameter	*Value*
a (semi-major axis)	6377563.396 m
b (semi-minor axis)	6356256.909 m
e^2	$(a^2 - b^2) / a^2$
F_0 (scale factor)	0.9996012717

Worked example: passing WGS84 coordinates (e.g. Wigan at lat 53.5, lng −2.6) into an EPSG:27700 endpoint expecting eastings/northings (359092, 400613) would return zero spatial features and silently contaminate the cohort with false negatives.

SYSTEMS ENGINEERING: TRI-STATE TYPING AND COMPUTATIONAL FRICTION

This chapter documents the systems engineering required to conduct national-scale spatial research against state digital infrastructure. The most dangerous trap in computational data science is the silent conversion of integration failures into negative observations: a timed-out query returning { inCoalArea: false } transforms a broken server into mathematical evidence of a hazard-free landscape. The chapter documents the tri-state enumeration architecture that abolishes binary boolean fallbacks.

Engineering layer	*Problem addressed*	*Solution*
Bounded concurrency	HTTP 429/503 firewall triggers	p-queue throttling; max 1–4 workers; 250ms jitter
Memory-safe lifecycle	Async timer leaks (V8 heap exhaustion)	Explicit clearTimeout() on resolve and reject
Tri-state typing	Error-as-false contamination	HazardStatus ∈ {TRUE, FALSE, ERROR, OUT_OF_BOUNDS}
Offline spatial grid	**NaPTAN — no live per-site API**	***In-memory 2D spatial grid over 400,000 DfT nodes (a design choice, not a measured failure)***

State digital infrastructure fragility:

Statutory dataset	*Managing authority*	*Observed result*	*Failure modality*
Coal Mining Risk Areas	Coal Authority / MHCLG	Configured URLs invalid; catalogue layer responds	Three configured URLs return HTTP 200 with JSON “Invalid URL” errors; a catalogue-discovered MRA layer returns a feature
Major Accident Hazard Pipelines	**Health & Safety Executive**	**No national service found**	Pipeline queries two reachable council layers; a national-service absence is not established by probe
HM Land Registry INSPIRE Index Polygons	HM Land Registry	Scripted 403; browser says WFS disabled	Layered refusal observed from one vantage point; cause not inferred

Audit note (Coal): the earlier “100% ERROR / no open vector REST API / commercial revenue protection” framing is withdrawn. The MRA public ArcGIS catalogue responds, and a catalogue-discovered DevelopmentHighRiskArea layer returns a feature at the test site. Three of four statically configured URLs return HTTP 200 with a JSON “Invalid URL” error; the fourth, under an organisation whose publisher is unknown, returns no features at that site. The cohort's coal results are 28 TRUE / 624 FALSE / 0 ERROR (two runs; neither recorded which address it used). Code still returns tri-state ERROR when no endpoint passes the benchmark — never coerced to “no coal risk.”

On the Land Registry finding: a public dataset (freehold title boundaries) served through a standard machine interface refused this research client's scripted requests at the network edge from one vantage point on one day; a browser request the same day was told the service was disabled. The response states no reason, and none is inferred here. The edge refusal may be generic automated-traffic protection, and INSPIRE obligations have not been assessed against these observations. The ownership analysis therefore uses bulk INSPIRE / Price Paid / CCOD releases — a workaround available to a resourced research pipeline, not necessarily to a community within a 21-day window.

STATUTORY DEFENSIBILITY — TWO AXES, NOT ONE

This chapter measures whether statutory planning protections scale with wealth, using a full-gradient non-parametric test. The Log-Defensibility Ratio, L(D,i) = ln(N_protective,i + 1) − ln(N_hazardous,i + 1), resolves the mathematical volatility of division-by-zero.

Full-gradient result:

IMD Decile Band	*n*	*Mean Protective Shields*	*Mean Hazard Burdens*
High Deprivation (1–3)	87	0.34	0.05 (n = 87)
Intermediate (4–7)	321	0.50	0.02 (n = 321)
Low Deprivation (8–10)	242	0.64	0.07 (n = 242)

Inferential test (Mann-Whitney U, deprived vs. affluent, flood removed): U = 8,370.5, z = −3.14, p = 0.0017, r = 0.17 (n = 87 / 242).

- **Ecological shields (Ancient Woodland, Green Belt, Conservation Area): affluent-skewed, flood-removed composite p = 0.0017**
- **Heritage at Risk: deprived-skewed, OR = 7.54, p = 3.2×10^{-5}**
- **Five-shield composite v5 p = 0.14 (N = 641); not yet recomputed on N = 652**

Extreme-groups comparison:

Statutory designation	High Dep. (n=87)	Low Dep. (n=242)	Disparity (pp)	OR (95% CI)	p-value	Status
Ancient Woodland Buffer (Shield)	9.2%	33.1%	+23.9	4.88 [2.25, 10.58]	8.4×10^{-6}	Significant
Flood Zone 2 detected (invalid exposure)	**17.2%**	**8.7%**	**+8.5**	**2.19 [1.07, 4.48]**	**0.043**	***Not a valid exposure measure****
Designated Green Belt (Shield)	***9.2%***	***16.9%***	***+7.7***	***2.01 [0.90, 4.49]***	***0.11***	***Inconclusive***
Conservation Area Designation	***14.9%***	***11.2%***	***+3.7***	***1.40 [0.69, 2.85]***	***0.35***	***Inconclusive***

Affluent areas show significantly higher odds of Ancient Woodland proximity. The flood layer is excluded from the composite because the pipeline wrote undeterminable values as “Zone 1 assumed”; flood data are not a valid exposure measure. The flood-removed composite is U = 8,370.5, z = −3.14, p = 0.0017, r = 0.17 (n = 87 / 242). Ancient Woodland is 9.2% (8/87) in high-deprivation sites versus 33.1% (80/242) in low-deprivation sites (OR = 4.88 [2.25, 10.58], Fisher p = 8.4×10^{-6}). Green Belt is 9.2% vs 16.9% (OR = 2.01 [0.90, 4.49], p = 0.11); Conservation Area designation is 14.9% vs 11.2% (OR = 1.40 [0.69, 2.85], p = 0.35). Heritage at Risk is the counter-axis: 16.1% (14/87) vs 2.5% (6/242), OR = 7.54 [2.80, 20.34], Fisher p = 3.2×10^{-5}. The five-shield composite has not been recomputed on N = 652; its v5 result was p = 0.14 (N = 641). These findings distinguish affluence-linked ecological shields from deprivation-linked built-heritage risk, not a single uniform measure of protection.

OPERATIONAL LANDSCAPES AND BMV FARMLAND

This chapter tests whether speculative infrastructure capital intersects high-grade agricultural land disproportionately compared with urban infill development. Sites were classified into four development archetypes by lexical parsing of campaign titles and planning application descriptions, then cross-referenced against Natural England's Provisional Agricultural Land Classification layer.

Development Archetype	*n*	*Provisional BMV (%)*	*Green Belt (%)*	*Mean IMD Decile*
Grid-Scale Energy Capital	55	72.7% (40/55)	10.9%	5.9
Logistics & Freight Depots	18	50.0% (9/18)	16.7%	6.6
Strategic Residential Sprawl	347	54.2% (188/347)	19.9%	6.9
Urban Infill & Estate Clearance	**232**	**46.6% (108/232)**	**0.4%**	**5.8**
Full Canonical Cohort	**652**	**52.9% (345/652)**	**12.1% (79)**	**6.4**

Inferential test: Pearson's χ^2 on the 4×2 contingency table: $\chi^2 = 12.72$, $df = 3$, $p = 0.0053$ (all 652 sites).

The ALC Grade 3a/3b distinction:

Category	*Definition*	*Cohort n*	*Cohort %*
Confirmed High-Grade BMV	Grade 1 and Grade 2 only	91	14.0%
Provisional-Inclusive BMV	Grades 1, 2, and undivided Grade 3	345	52.9%

Under NPPF Annex 2, ALC Grade 3b is not BMV farmland; Grade 3a is. Natural England's Provisional ALC does not subdivide Grade 3. The finding is a statistical association, not a causal demonstration. Level 2.

*Robustness: the BMV × archetype association survives whole-thesis Benjamini-Hochberg FDR (adjusted p = 0.0095) but not whole-thesis Bonferroni across the 18 valid tests (adjusted p = 0.095). Reported accordingly.

OPERATIONAL LANDSCAPES AND BMV FARMLAND

This chapter tests whether speculative infrastructure capital intersects high-grade agricultural land disproportionately compared with urban infill development. Sites were classified into four development archetypes by lexical parsing of campaign titles and planning application descriptions, then cross-referenced against Natural England's Provisional Agricultural Land Classification layer.

Phase-2 — Brownfield Register (planning.data.gov.uk): only 31 of 652 sites (4.8%) are on or adjacent to a registered brownfield parcel. The archetype cross-tab debunks the “infill” label: “Urban Infill & Estate Clearance” registers at just 6.9% (16/232) brownfield versus 3.6% (15/420) for every other archetype; full-table $\chi^2(3) = 5.52$, $p = 0.14$. 93.1% of “infill”-labelled sites are not on the register at all — the brownfield-reuse justification the label implies is not supported by the data. Independent corroboration: the government’s own digital-land platform team has publicly documented that council-published brownfield data was too inconsistent to aggregate reliably — government-sourced evidence for the same state-infrastructure-fragility argument made in Chapter 6. (Coordinates validated against an England bounding box to reject the platform’s own documented swapped-coordinate “in the sea” records.)

Phase-2 — Land Opacity, Mechanism A (HM Land Registry CCOD + Price Paid): of the 652 sites, 397 matched a corporate-owned title and 255 (39.1%) did not — 39.1% unresolved corporate ownership via postcode approximation. CCOD lists corporate proprietors only and every match is a postcode approximation, so this measures unresolved corporate ownership, not unregistered land (Level 1). In v5, of the 396 matched titles, 385 were currently corporate-owned; separately, 42 sites’ most-recent freehold seller matched a public-sector naming pattern (council, NHS, MOD, Forestry Commission) — these two counts describe different things (current ownership type vs. most-recent seller type) and are not mutually exclusive, so they are not meant to sum to 396. Mechanism B (option agreements / land banking) is explicitly not tested — no public register of option agreements currently exists, though a government proposal (associated with Planning Minister Matthew Pennycook) to create one is live policy. Named as future work; never inferred indirectly. Corporate ownership denominator (v5, n = 309, Track E1): the resolved corporate proprietors classify — via recursive Companies House PSC/BFS traversal — into STRICT (7 volume housebuilders: Barratt Redrow, Bellway, Crest Nicholson, Taylor Wimpey, MJ Gleeson; 2.3%), BROAD1 (24 developers, 7.8%), and BROAD2 (35 housebuilding-SIC holders, 11.3%). Developer ownership is now a measured variable (feeds the Track E3 test), superseding the earlier inconclusive spv_suspected = 0.

ECOLOGICAL DATA POVERTY AND THE PERI-URBAN VOID

This chapter computes baseline biological data availability across the canonical cohort and tests whether it tracks deprivation. Historical NBN Atlas species records were counted within a 1,000m radial buffer of each site. An "Ecological Data Desert" is defined as fewer than 5,000 historical records. The finding is counter-intuitive: data scarcity is non-monotonic, peaking in the intermediate band — the peri-urban fringe where speculative capital is most active.

The Flipped Archetype Finding (Track E2 & E3): Strategic Residential Sprawl is under-represented in data deserts: 28.5% (97/340) vs. 42.1% (126/299) for all other archetypes (OR = 0.55 [0.39, 0.76], Fisher's p = 0.00035). Data deserts concentrate instead in utility-scale energy (63.6%) and freight logistics depots (47.1%). Developer-owned sprawl link test (Track E3): OR = 0.42 [0.09, 1.95], Fisher's p = 0.35 (clean inconclusive null, provisional pending the E1 classification for 111 new v6 sites). The "developers target data deserts" hypothesis is refuted for residential development. Level 2 (statistical association).

Full-gradient result:

IMD Decile Band	*n*	*Deserts*	*Desert %*	*Mean Records (1km)*	*Mean Ec Score (0–100)*
High Deprivation (1–3)	86	30	34.9%	13,783	81.3
Intermediate (4–7)	317	122	38.5%	11,949	78.6
Low Deprivation (8–10)	234	70	29.9%	13,183	82.0
Full Canonical Cohort	**639**	**223**	**34.9%**	**12,644**	**80.2**

13 of 652 sites have no NBN count (the lookup failed) and are excluded, not counted as deserts.

Per-decile breakdown of the intermediate band:

IMD Decile	*n*	*Deserts*	*Desert %*
4	52	16	30.8%
5	72	27	37.5%
6	97	46	47.4%
7	96	33	34.4%

Inferential test (deprived vs. affluent, mean Epistemic Completeness): Mann-Whitney U = 9,429, z = −0.95, p = 0.34, r = 0.05.

There is no detectable class difference in baseline ecological data availability. The desert rate is non-monotonic and peaks in the intermediate band, driven by IoD2025 decile 6 (47.4%) — the commuter-belt fringe, neither urban enough to attract civic data-mapping nor ecologically designated enough to attract formal conservation surveys. The disparity here is financial, not class-targeted specifically: affluent groups can fund private ecological counter-surveys (£10,000–£30,000); disadvantaged communities generally cannot. Level 1 (desert rate); Level 2 (null class result).

COMPOUND EXPOSURE AND TRIPLE JEOPARDY

This chapter isolates frontline communities enduring simultaneous economic poverty, environmental burden, and active development conflict. The Triple Jeopardy threshold is defined as falling within the bottom 30% nationally for both economic deprivation (IMD 1–3) and environmental quality (IMED 1–3). Clinical health exceedances were cross-tabulated against the age-standardised IMD Health Deprivation and Disability Domain (File 7) to control for demographic aging.

Triple Jeopardy cohort:

Population Frame	*Triple Jeopardy Sites*	*Share of Frame*
National campaign harvest (N = 852)	54	6.3%
Canonical cohort (N = 652)	47	7.2%

Recomputed on the strict v6 cohort. Triple Jeopardy requires both the deprivation decile and IMED resolved, so it is defined only on the canonical cohort.

Compound exposure statistic:

Measure	*Triple Jeopardy subset (n=47)*	*Wider canonical cohort (N=652)*
Above England average for ≥1 condition	72.3% (34 sites)	-
Above avg for ≥1 of all 9 conditions (robustness)	72.3% (34 sites)	90.0%
Interpretation	Descriptive compound exposure	Not a compounding effect (cohort-wide health challenge)

Clinical health exceedances (QOF 2024–25; health-domain decile ≤ 5 vs > 5):

Clinical Condition	Deprived Half (% above avg)	Affluent Half (% above avg)	Disparity (pp)	OR (95% CI)	p-value
Obesity Prevalence	65.8%	34.4%	+31.4	3.67 [2.57, 5.22]	2.0×10^{-13}
COPD (Chronic Respiratory)	**67.9%**	**37.5%**	+30.4	**3.53 [2.47, 5.04]**	$\mathbf{1.1 \times 10^{-12}}$
Diabetes Mellitus	***64.8%***	***35.5%***	+29.3	***3.34 [2.35, 4.75]***	$\boldsymbol{9.6 \times 10^{-12}}$
Depression Prevalence	***67.9%***	***47.5%***	+20.4	***2.34 [1.64, 3.33]***	$\boldsymbol{2.0 \times 10^{-6}}$

Transit isolation alone is not a national driver of respiratory illness (OR = 0.80, p = 0.54). The compound exposure statistic describes a population, not a compounding effect: on the corrected join the Triple Jeopardy subset (72.3% on the all-9 measure) sits below the wider canonical cohort (90.0%). Level 1 (clinical exceedances); Level 2 (compound exposure association).

POLITICAL GEOGRAPHY OF CONTESTATION

This chapter characterises the political geography of contested planning campaigns relative to the post-July 2024 House of Commons composition. Campaigns over-index in Conservative and Liberal Democrat constituencies relative to seat share. Labour-held constituencies, however, account for the overwhelming majority of Triple Jeopardy battlegrounds, producing two distinct geographical dynamics.

Cross-party parliamentary exposure (N = 852 harvest frame):

Parliamentary Party	*Contested Sites*	*% Campaigns*	*Commons Seat Share*	*Disparity*
Labour / Co-operative	369	43.3%	63.2% (411)	−19.9 pp
Conservative	268	31.4%	18.6% (121)	+12.8 pp
Liberal Democrat	142	16.6%	11.1% (72)	+5.5 pp
Reform UK / Green / Others	73	8.7%	7.1% (46)	+1.6 pp

Two distinct geographical dynamics:

Dynamic	*Party concentration*	*Type of conflict*
Peri-Urban & Rural Friction	Conservative (+12.8 pp); Lib Dem (+5.5 pp)	Greenfield releases; solar/BESS; suburban expansion
Urban Compounding Core	Labour (87.2% of canonical Triple Jeopardy — 41/47; OR = 9.51 [3.98, 22.73], Fisher's exact $p = 7.4 \times 10^{-10}$; +45.4 pp vs 41.8% of non-TJ)	High-density infill; logistics depots

This is descriptive political geography, not a measurement of partisan motivation. Campaign volume concentrates in Conservative/Lib Dem seats, but compound vulnerability concentrates in Labour seats — the class-based injustice argument operates at a different spatial resolution than raw campaign counts. Level 1 (direct observation) for the campaign-volume geography; Level 2 (statistical association) for the Triple-Jeopardy–Labour concentration (OR = 9.51, Fisher's exact $p = 7.4 \times 10^{-10}$; robust under whole-thesis Bonferroni — Slide 19). Party is the sitting MP's party recorded at enrichment (Sept 2026); the earlier 34/52 figure had no generating script and is replaced. No motivational claims are made.

POLITICAL GEOGRAPHY OF CONTESTATION

Chapters 7–11 measure asymmetry at the input (consultation) side. The Phase-2 outcome layer opens the sharper question: does that asymmetry predict what actually happens to a site?

Cross-party parliamentary exposure (N = 852 harvest frame):

Source	*Granularity*	*Coverage*	*What it gives*
PINS (Planning Inspectorate quarterly stats)	LPA-level (aggregate)	205/206 cohort LPAs matched (190 OK, 15 LOW_N)	Appeal allow-rate per LPA — an ecological variable, not site-level; the live case-level dataset has been defunct since 2018
PlanIt (~420 council portals)	Site-level	650/652 sites resolved	Actual application state: Permitted / Conditions / Rejected / Withdrawn / Undecided

Source	Granularity	Coverage	What it gives

The open hypothesis: sites with a more asymmetric pre-application information environment — low ecological baseline, unresolved/opaque ownership, statutory-API failure — show systematically different outcomes, independent of the underlying constraint profile (BMV, Green Belt). Plainly: do developers win more often in Ecological Data Deserts?

Result: of 606 decided sites with a resolved NBN count, 546 permitted, 60 refused — a 90.1% developer success rate. Data-desert sites show no detectable difference in permission: OR = 0.98 [0.56, 1.71], Fisher p = 1.00; unresolved corporate ownership OR = 0.75 [0.44, 1.28], p = 0.34 (n = 619 decided sites). This is a clean null — and it supersedes the earlier 47%-coverage signal (OR = 2.45, p = 0.11), which was flagged as exploratory and did not survive powering-up. The honest reading: the consultation-stage asymmetry does not translate into a detectable difference in the binary permit/refuse outcome — which bounds the thesis to the consultation side and forecloses any outcome-determination claim it never made. Two caveats: an ~90% permit ceiling leaves little variance for any predictor to explain, and the binary flag can't see what gets built (scale, conditions). PlanIt's rate-limited backlog remains a civic-infrastructure fragility signal (Chapter 6). Level 1 for the matched records; Level 2 tested null for the prediction. No Level 4 claim.

WHOLE-STOCK SLACK VS. GREEN BELT RELEASE

Testing the "greenfield supply necessity" premise against the 10 GMCA boroughs' own published stock accounts. The tri-level continuum: micro (N = 652 sites) ↠ meso (whole stock vs. Green Belt) ↠ macro (Slide 14c).

The Green Belt slack ratio:

Measure	***Value***
Long-term empty dwellings (> 6 months)	12,824
15-year Places for Everyone Green Belt target	21,030 dwellings / 2,715 ha
Long-term empty as share of the GB target	4,967 = **166.3%** of the GB target
Boroughs where long-term-empty stock alone exceeds the 15-yr GB allocation	**Bolton (157.8%), Salford (207.8%)**

The 1% annual flow ceiling: new-build completions run at a 14-year mean of 0.54%–0.85% of standing stock (0.26%–0.76% in 2024/25). Pre-existing stock is >98% of physical supply in any year — the largest policy lever is the existing stock, not greenfield addition.

The exhausted-capacity premise for Green Belt release is not supported across GMCA. Level 1 (direct observation), from cryptographically-hashed gov.uk workbooks.

THE AFFORDABILITY WALL & YOUTH TENURE LOCKOUT

Macro tier: who the released supply is actually for.

- The 4.5× borrowing wall. ONS lower-quartile price-to-earnings ratios: 5.39× (Oldham) ↠ 8.81× (Trafford) — all 10 boroughs above the ~4.5× mortgage-multiple ceiling. Single lower-quartile buyers are priced out everywhere.
- Working-age PRS lockout. Private renting among working-age households: 37.6% (Manchester), 32.7% (Salford); working-age ownership collapsed to 34.7% and 42.8%; over-65 outright ownership 60%–77%.
- Pushed into ageing fabric. Young renters concentrate in pre-1919 Victorian/Edwardian stock — Oldham 35.4%, Rochdale 34.3%, Tameside 31.1%, Bolton 30.0%.

Tri-level synthesis: site-level epistemic enclosure (unresolved corporate ownership 39.1%, BMV 52.9%, 21-day compression) ↠ metropolitan asset management (Green Belt released while empties alone would cover most of the plan) ↠ a supply unaffordable to the younger cohort it is nominally for. Two faces of one distributional settlement. Level 1 for all figures; the synthesis is a labelled Level-3 framing.

THE CPSS ARCHITECTURE

A critical epistemological vulnerability in applied AI is the assumption that Large Language Models (LLMs) can act as deterministic statutory instruments. LLMs are probabilistic text generators. The CPSS architecture resolves this by structurally constraining the LLM as a lexical extraction engine operating entirely within a deterministic programmatic wrapper.

Before any objection letter is surfaced to a community, it must survive an adversarial panel of four validators:

Three architectural pillars:

Pillar	*Mechanism*	*Purpose*
Statutory Genotype	11 regenerative principles encoded as TypeScript classes	Translate regenerative ambition into statutorily compliant claims
Frame-Based Principled Reasoning	13 domain-specific Zod schemas	Force LLM to extract typed structured data (e.g. TransportAssessment, FloodRiskAssessment)
Four-Validator Manifold	4 independent validators + GATE-4 halt	Adversarial quality control before output surfaces to community

The Four-Validator Manifold

Validator	*Rule set*	*Failure mode caught*
Case-Law Factbook	8 canonical cases (e.g. R (Finch) [2024])	LLM hallucinated case citations
Forbidden-String Corpus	23 rules across 6 leak classes	Internal pipeline tags; "sales-style" AI language
Archetype × Ground Matrix	Statutory applicability rules	Non-applicable objections (e.g. Active Travel on unmanned BESS)
Citation Provenance Validator	December 2024 NPPF statutory allowlist	Invalid NPPF paragraph citations

If the aggregate forensic score falls below 40/100, the system triggers a GATE-4 Statutory Halt — structurally refusing to produce a non-viable output and generating a missing-documents schedule instead. The system measures within a procedurally defined space; it does not decide what is true.

THE CPSS ARCHITECTURE

A critical epistemological vulnerability in applied AI is the assumption that Large Language Models (LLMs) can act as deterministic statutory instruments. LLMs are probabilistic text generators. The CPSS architecture resolves this by structurally constraining the LLM as a lexical extraction engine operating entirely within a deterministic programmatic wrapper.

Before any objection letter is surfaced to a community, it must survive an adversarial panel of four validators:

Three architectural pillars:

Pillar	*Mechanism*	*Purpose*
Statutory Genotype	11 regenerative principles encoded as TypeScript classes	Translate regenerative ambition into statutorily compliant claims
Frame-Based Principled Reasoning	13 domain-specific Zod schemas	Force LLM to extract typed structured data (e.g. TransportAssessment, FloodRiskAssessment)
Four-Validator Manifold	4 independent validators + GATE-4 halt	Adversarial quality control before output surfaces to community

The Four-Validator Manifold

Validator	*Rule set*	*Failure mode caught*
Case-Law Factbook	8 canonical cases (e.g. R (Finch) [2024])	LLM hallucinated case citations
Forbidden-String Corpus	23 rules across 6 leak classes	Internal pipeline tags; "sales-style" AI language
Archetype × Ground Matrix	Statutory applicability rules	Non-applicable objections (e.g. Active Travel on unmanned BESS)
Citation Provenance Validator	December 2024 NPPF statutory allowlist	Invalid NPPF paragraph citations

If the aggregate forensic score falls below 40/100, the system triggers a GATE-4 Statutory Halt — structurally refusing to produce a non-viable output and generating a missing-documents schedule instead. The system measures within a procedurally defined space; it does not decide what is true.

THE 11 REGENERATIVE PRINCIPLES AS PLANNING AXIOMS

This chapter introduces the Community Planning Support System (CPSS) — the operational instrument built to close the evidentiary asymmetry documented in Chapters 7–11. CPSS is not a generic AI letter-writer. It is a statutory measurement instrument designed to measure within a procedurally defined space, ensuring that the AI measures within a space whose boundaries are procedurally defined rather than deciding what is true. The architecture rests on three pillars.

Principle	*Theoretical Source*	*Statutory Operationalisation*
Wholeness	Mang & Haggard (2016)	Requires all four capitals (financial, social, natural, human) to be represented in the capital map.
Nestedness	Lyle (1994); Reed (2007)	Forces evaluation at neighbourhood, regional, and planetary scales (Net Zero contribution).
Place	Lefebvre (1991)	Checks for epistemic plurality and historical depth to counter generic developer typologies.
Reciprocity	Indigenous Governance	Checks for "gives" and "receives" terms to counter transactional S106 framing.
Nodal Interventions	Donella Meadows (1999)	Binds the archetype × ground matrix to identify the highest statutory leverage points.
Pattern Literacy	Christopher Alexander (1977)	Evaluates counter-pattern identification and place-specific pattern generation.
Levels of Work	Elliott Jaques (1989)	Requires 75% coverage across operate, maintain, improve, and regenerate levels.
Temporal Justice	Wales Future Generations Act	Enforces a 7-generation horizon; invokes R (Finch) Scope 3 emissions and Climate Change Act 2008.
Distributed Authority	Elinor Ostrom (1990)	Computes dual Gini coefficients; fires hard violations if power Gini exceeds 0.45.
Generative Conflict	Frankfurt School	Bans "trade-off" and "compromise" language; forces constitutional adjudication.
Implementation Resistance	Anti-Capture Design	Checks transition pathways across legal, economic, cultural, and political categories.

CPSS is not a generic AI letter-writer — it is a statutory measurement instrument, designed to measure within a procedurally defined space rather than decide what is true.

THE NPPF 11(D) WORKED EXAMPLE

This chapter demonstrates the CPSS operating end-to-end on a specific statutory problem: the NPPF Paragraph 11(d) "tilted balance." When a Local Authority cannot demonstrate a 5-year housing land supply, Paragraph 11(d) forces planning officers to apply a presumption in favour of sustainable development, suspending local democratic plans. The CPSS engine operationalises Footnote 7 of Paragraph 11(d)(i), which legally disarms the tilted balance when specific protected assets are present.

The Unified Tetrad:

Stage	*Content*
THEORY	Capital seeks the "Spatial Fix" via speculative unallocated greenfields (Christophers)
EMPIRICAL PROBLEM	Peri-urban campaigns face developer claims under failed 5-Year Housing Land Supply
STATUTORY REALITY	NPPF Paragraph 11(d) forces the "tilted balance" in favour of development
STATUTORY INSTRUMENT	CPSS operationalises Footnote 7; scans 64 spatial layers; disarms the tilted balance

Footnote 7 assets scanned by CPSS:

Asset category	*Examples*
Ancient Woodland	Statutory ancient woodland inventory
Peatland	Deep peat designations
Flood Zone 3	Environment Agency flood mapping
Asset category	Examples

This worked example demonstrates that the pipeline is not merely diagnostic. It produces statutorily valid community representations that restore procedural parity by leveraging the state's own planning framework.

CONTRIBUTION TO KNOWLEDGE

This chapter synthesises the fourfold contribution of the research: methodological, empirical, theoretical, and applied. The thesis extends Tseng, Becker & Roikonen's (2024) "epistemic enclosures" concept from digital participatory-budgeting platforms into statutory and environmental evidentiary enclosure, and complements rather than supersedes Schmid & Topalović's (2023) qualitative mapping of operational landscapes.

Contribution type	*Content*
Methodological	Reproducible, tri-state-typed computational pipeline auditing 652 sites against 65 statutory layers
Empirical	Four inferential findings: Defensibility Gap (flood-removed composite p = 0.0017); BMV × archetype (p = 0.0053); ecological data scarcity (34.9%, class-test p = 0.34); state API fragility (three distinct failure modes)
Theoretical	Four inferential findings: Defensibility Gap — ecological-shield axis (Ancient Woodland $p = 8.4 \times 10^{-6}$; composite with flood removed p = 0.0017), with a deprivation-skewed heritage counter-axis ($p = 3.2 \times 10^{-5}$); BMV × archetype (p = 0.0053); universal ecological data scarcity (no class difference, p = 0.34); state API fragility — three services, three failure modes, one vantage point, 4 Oct 2026 (HM Land Registry INSPIRE: edge refusal, and Service WFS is disabled in a browser capture; HSE MAHP: no national service found; Coal/MRA: configured addresses invalid, layer discoverable through the catalogue)
Phase-2 empirical extension	Six linked datasets extend the audit beyond consultation: revealed car dependency (76.6%); Heritage at Risk as a 5th shield (45 of 652 sites, deprivation-skewed); brownfield reuse shown to be the exception rather than the rule (31/652); unresolved corporate ownership via postcode approximation (255 of 652 sites, 39.1%); PINS + PlanIt outcome layer answering the asymmetry-predicts-outcome question — a tested null (desert↛permit OR = 0.98, p = 1.00), bounding the thesis to the consultation side
Theoretical	Bounded evidence-ladder-compliant claim: informational architecture produces evidentiary asymmetry during 21-day consultation window (Epistemic Enclosure)
Applied	CPSS architecture: statutory measurement instrument operationalising NPPF 11(d)(i) Footnote 7 with GATE-4 halt

Novelty bounding:

Claim	*Scope*
Domain extension of Tseng et al. (2024)	From platform filtering to baseline evidence generation
Complementary audit	Not superseding Schmid & Topalović (2023)
First of its kind	First to subject statutory/informational preconditions to inferential statistical analysis at this scale

The contribution is explicitly bounded. The theoretical claim is a domain extension, not a coinage. The empirical finding is a statistical association, not a causal demonstration. The methodological innovation is a reproducible pipeline, not an unprecedented computational task.

LIMITATIONS AND OUTSTANDING VERIFICATION

This chapter documents the four methodological limitations governing interpretation, and the four outstanding verification items that should be resolved before final submission. It also restates the four-level evidence ladder applied throughout the thesis, under which no Level 4 (bounded causal assertion) claims are made.

Methodological limitations:

#	*Limitation*	*Mitigation*
1	Convenience sample; requires campaign organising capacity	Conservative bias direction stated explicitly
2	Health-domain join uses IMD 2019 (File 7)	IoD2025 deprivation deciles verified; IMD 2019 retained only for health-domain adjustment
3	Extreme-groups design excludes 49.2% intermediate cohort	Paired with full-gradient statistics
4	The MCAR assumption — excluding ERROR rows assumes API failures are unrelated to site characteristics	**Addressed** — Global Moran's I run: inferential layers (IMD/IMED/env) 0% missing (MCAR not engaged); flood-zone availability locally clustered, treated as suggestive

Outstanding verification items:

#	*Item*	*Status*
1	Independent GIS spot-check of 5–10 raw site intersections	Recommended before final submission
2	Done — IoD2025 deprivation deciles verified	Sensitivity test if available
3	Confirm ε-invariance result was re-run	Confirm from script output, not summary text
4	Compound exposure recomputed on N = 47 canonical subset	Done (v6): 72.3% (34/47); 90.0% cohort-wide

Evidence-ladder protocol:

Level	Definition	Example
1	Direct observation	"34.9% of 639 sites with resolved NBN counts intersect areas with fewer than 5,000 historical biological records"
2	**Statistical association**	**"Speculative energy capital exhibits a significant association with BMV farmland (χ^2 = 12.72, p = 0.0053)"**
3	***Plausible mechanism (must be labelled hypothesis)***	***"21-day windows combined with data scarcity create conditions under which applicants may assert low-baseline values"***
4	***Bounded causal assertion***	***Not made anywhere in this thesis***

The evidence-ladder protocol applies a canonical language sheet throughout. Banned terms include: "scales monotonically," "clinical compounding," "proves," "weaponise," "manufacture tabula rasa," "constitutional instrument," "coinage," and "unprecedented."

CONCLUSION

The empirical baseline establishes that statutory protections, baseline ecological data, and procedural rights are distributed unequally across England's planning system. These findings indicate that an asymmetry of technical and evidentiary capacity operates as a significant, measurable barrier to spatial justice — this thesis does not compare the relative size of this barrier against other candidate barriers (funding, legal representation, political influence), which lies beyond its scope.

Four conclusions:

#	*Conclusion*	*Key statistic*
1	Statutory Defensibility Disparities	Ecological/natural shields skew affluent (Ancient Woodland OR = 4.88, $p = 8.4 \times 10^{-6}$; composite with flood removed p = 0.0017); built-heritage protection skews deprived (Heritage at Risk OR = 7.54, $p = 3.2 \times 10^{-5}$) — not a contradiction, but a sharper account of which kind of protection tracks wealth
2	Peri-Urban Evidentiary Void	38.5% intermediate-band data deserts; IoD2025 decile 6 peak 47.4%; no class difference overall (p = 0.34)
3	Infrastructural Transformation of Farmland	Energy 72.7%; sprawl 54.2%; logistics 50.0%; infill 46.6%; $\chi^2 = 12.72$, p = 0.0053
4	State Digital Fragility	HM Land Registry INSPIRE: layered refusal (Cloudflare 403 to scripted requests from one vantage point; browser capture Service WFS is disabled); HSE MAHP: no national service found (not testable by probe); Coal/MRA: configured addresses invalid, layer discoverable through the catalogue — open in law, uneven in programmatic practice

Thesis statement:

Element	*Content*
Central finding	An asymmetry of technical and evidentiary capacity is a significant, measurable barrier to spatial justice in the English statutory planning system
Mechanism	Epistemic Enclosure, operating during the 21-day statutory consultation window
Policy direction the findings support	Open, federated Civic Spatial Data Infrastructure capable of translating lived environmental realities into verifiable planning evidence on more equal procedural terms — a direction the findings are consistent with, not one this thesis's design establishes as necessary or sufficient on its own
Theoretical position	A complementary computational-statutory audit, extending rather than superseding the existing literature

Carlos Arleo



github
ARIANDE / Published Research / RoRi / System Architectures / YPV / DNI 4.0 /Aitiopoietic Methodology
+447881647716 / c.arleo@localis-ai.uk