The Bartlett School of Architecture

BARC0141: Built Environment Dissertation

"Scoring the Sky: Zoning, Daylight, and the Design Spaces of Midtown Manhattan."

by

José Ignacio Kempff Seleme

Date: 01 / 09 / 2026

Word Count: 9 999

Dissertation submitted in part fulfilment of the Degree of Master

MSc Architectural Computation
The Bartlett School of Architecture
University College London

Abstract

For sixty-six years, Midtown Manhattan's height and setback rules were written around a limited resource: the sky visible from the street. In 1982, the prescriptive envelopes of 1916 and 1961 were replaced by a performance score, the Daylight Evaluation Chart (DEC), allowing any massing scoring highly enough from fixed vantage points. This thesis treats the regulation as the experimental variable and examines whether that change preserved the metric while altering the street-level sky.

A multi-objective evolutionary algorithm generates building massing under four regimes: (1) unregulated pre-1916, (2) the 1916 setback rule, (3) the 1961 tower-on-base, (4) the 1982 chart, across four lot typologies. The sixteen cells map the design space each regime permits. A retrospective evaluation scores fifteen as-of-right Midtown buildings, five per regime, using the same validated implementation.

The 1982 regime allows a materially wider design space than either prescriptive predecessor. The ordering pre-1916 > 1982 > 1916 > 1961 holds in every cell. Yet the widening does not decouple score from street. The off-front rank correlation between DEC score and centreline sky view factor never falls below +0.868.

The retrospective evaluation falsifies the regulation's calibration in both directions. The 1916 cohort scores 18 points below the 75% benchmark and 14 below the implementation's anchor. The 1961 stock scores 12 and 16 above. Across the cohort, the score–sky coupling attenuates to +0.53, or +0.56 across frontages, short of the strong, consistent coupling in both evaluations that falsification requires. Only the weaker limb of the divergence hypothesis survives, with conservative departures.

The contribution is threefold: zoning regimes read as generators of design space; an open, worked-example-anchored implementation of the DEC that declares the counting convention the chart leaves unstated, paired with the street-level sky measure it cannot take; and the realised Midtown stock scored against its governing metric.

Keywords

Daylight Evaluation Chart, performance-based zoning, sky view factor, building massing, multi-objective evolutionary optimisation, NSGA-II, design space exploration, hypervolume indicator, regulation as code, New York City.

Table of Contents

Abstract 3
Table of Contents 5
List of Figures 9
List of Tables 10
List of Equations 10
Glossary of Terms 11
1. Introduction 13
  1.1 The zoning envelope as a form-giver 13
  1.2 Daylight as the organising concern of Midtown zoning 14
  1.3 The calibration and the gap it leaves 14
  1.4 Hypotheses 15
    H1: Calibration 15
    H2: Envelope 16
    H3: Experience 16
2. Literature Review 17
  2.1 The 1916 Resolution and the setback principle 17
  2.2 The 1961 Resolution and the tower-on-base 17
  2.3 The 1982 Resolution and the turn to performance 18
  2.4 The DEC as a measuring instrument 19
  2.5 The regulation since 1982 21
  2.6 SVF and street-level openness 21
  2.7 Computational and evolutionary approaches to building form 22
3. Methodology 23
  3.1 Overall approach: two evaluations 24
  3.2 Representing the regimes as constraints 24
  3.3 The Daylight Evaluation fitness function 25
  3.4 The evolutionary algorithm 26
  3.5 The SVF as an independent measure 27
  3.6 Experimental protocol and the confounding problem 27
  3.7 Genotype, decoding, and the treatment of legality 28
  3.8 Retrospective evaluation of the built record 29
  3.9 The instruments 31
4. Results 37
  4.1 H1 — calibration: cohort DEC scores 37
  4.2 H2 — envelope: HV across regimes 38
  4.3 H3 — experience: the DEC–SVF relationship 44

4.4 Regime findings: the building stock by regime 48
5. Discussion and Conclusion 53
5.1 Interpretation 53
5.2 The confound revisited 53
5.3 Limitations 54
5.4 Contribution and future work 56
References 59
Works cited 59
Primary regulatory sources 60
Historical regulatory instruments 61
Software and data 62
Appendices 63
Appendix A: Regulatory Regime Reconstructions 63
A.1 Pre-1916 baseline 63
A.2 1916 setback regime 63
A.3 1961 tower-on-base regime 64
A.4 1982 DEC regime 66
A.5 Scoped-out provisions 66
A.6 Encoding the regimes as constraint modules 66
A.7 Objective framing and the placement of the legality penalty 67
Appendix B: The Retrospective Cohort 68
B.1 MAPPLUTO screening criteria 68
B.2 The fifteen buildings 68
B.3 Setback verification (1916 cohort) 69
B.4 Regime cohort sizes 70
B.5 Frontage parameters and the vantage construction 70
B.6 Cohort scores 71
B.7 Per-frontage record 72
B.8 Compliance routes and subdistricts 73
B.9 Screening annex: the 1982 replacement 73
Appendix C: Instrument Validation 75
C.1 DEC scorer validation against the worked example 75
C.2 SVF sampling specification 76
C.3 HV computation and validation 77
C.4 The published Daylight Evaluation Diagrams 77
C.5 The scorer algorithm 79
C.6 The profile curve and the aggregation of scores 80
C.7 Component decomposition of the worked example 80

C.8 Sampling protocol for the coupling test 81
C.9 Convention-band sensitivity of the verdicts 81
Appendix D: Evolved Design Spaces in Full 83
D.1 Design-space envelopes by regime 83
D.2 Convergence and coupling 83
D.3 Design spaces by regime 85
D.4 Full numerical results 90
Appendix E: Component Structure 92
E.1 PhenotypeBuilder: Genome to massing 93
E.2 RegimePre1916: The unconstrained baseline 95
E.3 Regime1916: Height districts and the setback envelope 95
E.4 Regime1961: The 1961 prescriptive rule gates 97
E.5 Regime1982: Legality as daylight performance 99
E.6 DECProjection: The rectilinear chart projection 100
E.7 DECScorer: The production scorer 102
E.8 CenterlineSVF: The independent openness sampler 104
E.9 HypervolumeCalc: Measuring the design spaces 105

List of Figures

Figure 1.1 — Hugh Ferriss, Evolution of the Set-back Building, 1922 13
Figure 2.1 — The three regimes as envelopes drawn by the regulation 18
Figure 2.2 — The Daylight Evaluation Chart 19
Figure 2.3 — The SOM Rectilinear Chart: the rectilinear successor 20
Figure 3.1 — The four canonical lot typologies of the evolutionary evaluation 23
Figure 3.2 — Along-frontage sky-view profile of the worked-example massing 26
Figure 3.3 — The retrospective cohort as modelled 30
Figure 3.4 — The worked example: Kwartler and Masters (1984, Fig. 10) 32
Figure 3.5 — The worked-example massing, scored by this study's instrument 33
Figure 4.1 — DEC scores of the retrospective cohort, by regime 37
Figure 4.2 — HV ratios by regime and typology, as grouped bars 39
Figure 4.3 — HV ratio across the four typologies, one line per regime 40
Figure 4.4 — Pareto envelopes by regime for each of the four typologies 41
Figure 4.5 — One representative individual per cell, I 42
Figure 4.6 — One representative individual per cell, II 43
Figure 4.7 — DEC score against centreline SVF, pooled by typology 46
Figure 4.8 — The retrospective cohort 47
Figure 4.9 — The retrospective cohort in its screened population 49
Figure 4.10 — The Chanin Building (1916 cohort) at its binding vantage 50
Figure 4.11 — 880 Third Avenue (1961 cohort) at its binding vantage 51
Figure 4.12 — 425 Lexington Avenue (1982 cohort) at its binding vantage 52
Figure A.1 — The 1961 type at its two floor-area tiers 65
Figure C.1 — Daylight Evaluation Diagram: 60-foot street 78
Figure C.2 — Daylight Evaluation Diagram: 80-foot street 78
Figure C.3 — Daylight Evaluation Diagram: 100-foot street 79
Figure D.1 — HV convergence by regime, overlaid within each typology 84
Figure D.2 — DEC–SVF rank correlation, all sixteen cells 85
Figure D.3 — Pre-1916: evolved population (light) and Pareto front by typology 86
Figure D.4 — 1916: evolved population (light) and Pareto front by typology 87
Figure D.5 — 1961: evolved population (light) and Pareto front by typology 88
Figure D.6 — 1982: evolved population (light) and Pareto front by typology 89
Figure E.1 — PhenotypeBuilder 93
Figure E.2 — RegimePre1916 95
Figure E.3 — Regime1916 96
Figure E.4 — Regime1961 97
Figure E.5 — Regime1982 99

Figure E.6 — DECProjection 100
Figure E.7 — DECScorer 102
Figure E.8 — CenterlineSVF 104
Figure E.9 — HypervolumeCalc 105

List of Tables

Table 3.1 — DEC scorer validated against the ZR worked example 34
Table 4.1 — Pareto-front HV by regime and lot typology 38
Table 4.2 — Off-front DEC–SVF Spearman correlation, the load-bearing H3 test 44
Table 4.3 — On-front Spearman correlation, the secondary check 44
Table B.1 — The retrospective cohort: fifteen as-of-right Midtown buildings 69
Table B.2 — As-of-right office stock in Midtown by regime 70
Table B.3 — Retrospective cohort scores 71
Table B.4 — Per-frontage scores and sky measurements 72
Table B.5 — The 1982 replacement screen 74
Table C.1 — Verdict sensitivity across the blockage-convention band 82
Table D.1 — Raw HV by regime and typology 90
Table D.2 — Off-front integrity gate: primary sample size per cell 90
Table D.3 — Per-cell summary: coupling, gate exclusions and convergence 91

List of Equations

(E.1) 101
(E.2) 101

Glossary of Terms

BBL — Borough–Block–Lot. New York City's unique tax-lot identifier, recorded in MAPPLUTO; the key by which the retrospective cohort is identified (Table B.1).

CD — Community District. New York City's administrative subdivision; the screen of B.1 covers the three Midtown districts, coded 104–106 in MAPPLUTO.

DCP — Department of City Planning. The New York City agency that maintains the Zoning Resolution and publishes MAPPLUTO, the tax-lot dataset from which the retrospective cohort was screened (Appendix B.1).

DEC — Daylight Evaluation Chart. The performance instrument of the 1982 Midtown zoning (ZR§§81-271 to 81-275): building massing is projected onto a gridded chart of the sky from fixed vantage points on the street centreline, and scored by the sky it leaves visible. The thesis's central object; implementation validated at §3.9.

FAR — Floor-area-ratio. Built floor area divided by lot area, the density currency of the 1961 Resolution; the 1916 regime had no floor-area ceiling.

FAS — Fifth Avenue Subdistrict. A subdistrict of the Special Midtown District along Fifth Avenue (ZR §81-253), carrying its own street-wall and bulk controls that modify the standard daylight evaluation for lots within it. Two 1982 cohort buildings, 505 and 563 Fifth Avenue, lie inside it; they are retained in the cohort under a recorded flag rather than excluded, and were not held to the unmodified chart (Figure 4.1; B.8).

HV — Hypervolume. The measure of a Pareto front's dominated objective-space volume (Guerreiro, Fonseca and Paquete, 2021); the design-space currency of H2, reported raw in Table D.1 and as ratios to the pre-1916 baseline in Table 4.1.

MAPPLUTO — The New York City Department of City Planning's tax-lot dataset (Primary Land Use Tax Lot Output, joined to lot geometries): one record per tax lot carrying use class, dimensions, floor area, year built and the BBL key. The source of the retrospective screen and cohort (Appendix B.1); release as cited in the reference list.

NSGA-II — Non-dominated Sorting Genetic Algorithm II (Deb et al., 2002). The multi-objective evolutionary algorithm, as implemented in Wallacei, that evolves each regime's daylight–volume Pareto front in the evolutionary evaluation.

NYC — New York City.

OLS — Ordinary least squares. The linear regression fitted through the score–sky observations in Figures 4.7 and 4.8; descriptive only, the reported statistics being rank correlations.

SOM — Skidmore, Owings & Merrill. The architectural practice whose rectilinear redrawing of the DEC (Figure 2.3) supplies the projection geometry this implementation adopts (Equations E.1–E.2).

SVF — Sky View Factor. The fraction of the sky hemisphere visible from a point, cosine-weighted; sampled here along the street centreline at pedestrian height as the independent measure of street-level openness (§3.5).

ZD — Zoning District. The lot's zoning designation in the tax-lot record (e.g. C5-3).

ZR — Zoning Resolution. The New York City Zoning Resolution; section references (e.g. §81-274) are to its text as amended.

1. Introduction

1.1 The zoning envelope as a form-giver

In 1922, the architectural drafter Hugh Ferriss, working with the architect and zoning advocate Harvey Wiley Corbett, produced four charcoal drawings that would become among the most influential architectural images of the decade. Titled the *Evolution of the Set-back Building*, they did not depict a building. They depicted a regulation. The first drawing showed the raw maximum mass that New York's 1916 Zoning Resolution (ZR) permitted on a block: a vast sloping solid whose receding planes were nothing more than the angle of light the law required to keep sunlight on the street. The following drawings carved that mass into something orthogonal and buildable, step by step, until a recognisable stepped-pyramid skyscraper emerged. The point of the series was radical for its moment: the shape of the building was latent in the law before any architect touched it. Ferriss had not represented a finished design; he had manipulated building form by drawing directly from the zoning envelope, treating the regulation as the generative instrument (Ferriss, 1929).

The drawings were widely exhibited and published, and their effect on a generation of architects was to reframe the 1916 law not as a restriction but as a form-giving principle. This is the founding intuition of the present thesis: in Midtown Manhattan the regulation draws the building. The stepped towers of the 1920s and 1930s, the plaza-fronted slabs of the 1960s, the daylight-scored forms of the 1980s: each is the visible signature of a particular regulatory regime.



Figure 1.1 — Hugh Ferriss, Evolution of the Set-back Building, 1922. Source: Avery Architectural & Fine Arts Library, Columbia University.

1.2 Daylight as the organising concern of Midtown zoning

Zoning regulations distribute a limited resource. In the dense commercial city that resource is access to the sky, namely to daylight, to air, and to the sense of openness that distinguishes a street from a tunnel. New York City's 1916 ZR, the first comprehensive ordinance of its kind in the United States, was enacted to arrest the loss of that resource, its immediate provocation the Equitable Building of 1915. The resolution that followed made daylight the governing variable of tall-building form, and every major revision to Midtown's height and bulk rules since has renegotiated the same concern under changing economic and political pressure.

The 1916 Resolution enforced a setback envelope; the 1961 omnibus amendment replaced it with a floor-area-ratio (FAR) system, trading the stepped wedding-cake for the tower-on-a-plaza; and in 1982 the Midtown ZR abandoned prescription for performance, scoring each proposal by how much of the sky remained visible from the street. Chapter 2 describes the three regimes in detail.

Across sixty-six years and three regulatory philosophies, the constant was the claim that controlling the sky visible from Midtown streets produces streets worth being on. Daylight is never quite regulated for its own sake: it stands in for the openness and legibility of pedestrian space, the sense of the city's vertical grain from the ground. The regulatory tradition assumes that what the rule measures and what the pedestrian experiences are aligned closely enough that protecting the first protects the second.

1.3 The calibration and the gap it leaves

That assumption became explicit, and measurable, in 1982. The DEC at the centre of the Midtown Resolution was not designed in the abstract. Kwartler and Masters (1984), who developed it for the Department of City Planning, derived its passing thresholds from the daylight performance of buildings already standing under the 1916 and 1961 rules. The required score of 75% overall, with no street frontage below 66%, was a measured value: the average daylight the earlier prescriptive regimes had in fact delivered, with the lower bound set by the worst acceptable case, an interior-lot building in a two-times height district under the 1916 rules. The performance system was deliberately calibrated to reproduce the daylight outcomes of the regimes it replaced while freeing architects from their envelope.

That calibration is a pair of empirical claims, and they have different standings. The first is a claim about the built record: that the thresholds match the measured performance of the as-of-right building stock. The regulation states it in its own text, and as far as this review can find, no one has tested it computationally against the buildings it describes. The second is a claim about the street: that a building which earns the score gives pedestrians the experience the score was written to protect. No documentation was found that tests it against the score. The Chart scores the sky from fixed regulatory vantage points on the street centreline, a bounded sweep weighted above seventy degrees; a pedestrian sees the street continuously, the whole hemisphere available, weighted by nothing but the cosine of incidence (§2.6). The two geometries are related but not equivalent, and nothing in the calibration guarantees that a building engineered to satisfy the first delivers the openness implied by the second.

This is the gap the thesis investigates. The central claim under examination is that the 1982 regime preserved its metric while losing the experience the metric was a proxy for: that by releasing the prescriptive envelope while holding the daylight score fixed, it opened a region of viable, fully compliant building form that delivers measurably worse sky openness at street level than the older prescriptive regimes ever permitted.

1.4 Hypotheses

The claim decomposes into three falsifiable hypotheses, tested across two complementary empirical evaluations: a controlled computational experiment that maps what each regime *permits*, and a retrospective evaluation of a representative sample of as-built tall buildings that measures what each regime *produced*.

H1: Calibration

The DEC's stated calibration holds on the built record: as-of-right Midtown buildings constructed under the 1916 and 1961 regimes score approximately 75% on the Chart on average, and the worst-case interior-lot building under the 1916 rules scores approximately 66%. This is the regulation's own empirical claim, examined by the retrospective evaluation alone, and falsified if the cohort scores fall materially below the threshold band once the documented measurement conventions of the implementation are accounted for (§3.8). Either outcome is a finding: if the calibration fails, the regime preserved a nominal metric mis-calibrated against its own benchmark.

H2: Envelope

At equivalent daylight performance, the 1982 regime admits a materially wider space of viable building form than the 1916 and 1961 regimes, measured as the hypervolume (HV) of the DEC score–volume Pareto front evolved under each regime on each lot typology. It is examined by the evolutionary evaluation, and falsified if the 1982 fronts do not dominate materially greater HV at matched daylight, in which case the formal freedom the 1982 regime advertised was nominal rather than substantive.

H3: Experience

Equivalent daylight scores do not imply equivalent street-level sky openness: the correlation between the DEC score and an independent geometric measure of pedestrian openness, the Sky View Factor (SVF) sampled along the street centreline, is weak or inconsistent, and specifically weaker within the region of form the 1982 regime opened than among forms admissible under the prescriptive regimes. It is examined by both evaluations: across the evolved populations, under the anti-circularity protocol of §3.6, and across the real building cohorts, where no selection pressure contaminates the relationship. It is falsified by a strong and consistent correlation in both evaluations, in which case the calibration assumption is justified, and the metric genuinely stands in for the experience.

The hypotheses are partly ordered. H2 is a premise of H3: if the 1982 regime opened no new region of form, there is no widened envelope in which experience could degrade. H1, by contrast, is independent of both: the calibration can fail while the envelope widened and the experience diverged, and that conjunction would itself be a coherent and stronger finding.

2. Literature Review

This chapter traces the three regimes, explains the Chart as a measuring instrument and records its present legal status, introduces the SVF, and reviews the computational research the thesis draws on.

2.1 The 1916 Resolution and the setback principle

The New York City ZR of 1916 set the basic geometry of Midtown's built form for the following four decades. It introduced height districts, and a system of mandatory setbacks tied to street width: a building could rise to a street-wall height proportional to the width of the adjacent street, after which its mass had to step back beneath a sky exposure plane of prescribed angle, with a tower of unlimited height permitted over a quarter of the lot (Willis, 1995). The plane's geometry was set to keep a defined cone of sky visible from the street. The resulting typology, namely continuous street wall, stepped middle, slender crown, became the dominant image of pre-war Midtown, given definitive expression in the Ferriss drawings and canonical realisation in the Empire State Building. This envelope, street-wall height, sky exposure plane and tower exemption, is what §3.2 encodes as the 1916 constraint module.

2.2 The 1961 Resolution and the tower-on-base

The 1961 omnibus amendment replaced the setback regime with a floor-area-ratio system. Bulk was now governed by a maximum ratio of floor area to lot area rather than a prescriptive envelope, and the new tower-on-base typology, a slender tower rising from a reduced footprint, the remainder of the lot left as open space, took Mies van der Rohe's Seagram Building as its admired prototype (Willis, 1995). Kwartler and Masters (1984) record that the 1961 amendments were intended to deliver daylight equivalent to the 1916 rules through different formal means: the tower over an open plaza was expected to leave as much sky visible from the street as the stepped wedding-cake, while affording more efficient floor plates. Whether that equivalence held in practice was contested at the time and was never tested against an independent geometric measure of street-level openness, a point to which this thesis returns. The concern predates the instruments: the form of the tower drew its most influential critique from Jacobs (1961), published the year of the amendment, on grounds not of daylight but of the life of the street, the first sustained argument that the ground-level pedestrian experience, rather than conformity to a planning standard, is the proper measure of such a form.

2.3 The 1982 Resolution and the turn to performance

By the late 1970s a generation of large towers approved through the discretionary review provisions of the 1961 resolution had produced daylight conditions measurably worse than those of buildings built as of right under either earlier regime. The mechanism, as Kwartler and Masters (1984) describe it, was the waiver of prescriptive controls in favour of negotiated design review: a process meant to raise architectural quality had instead produced buildings performing poorly on the very measure the zoning existed to protect.

The 1982 Midtown ZR answered with a performance-based alternative to the prescriptive daylight-compensation rules. Its centrepiece was the DEC, a modified Waldram diagram adapted to Midtown street geometry by Kwartler and Masters with the building physicist Harvey Bryan. In place of a dictated envelope, the Chart offered a score, and any massing that met the threshold was permissible. The passing standards (75% overall, 66% for any single frontage) were measured from the existing as-of-right stock, as §1.3 recorded; the calibration was deliberate; it is the historical fact on which the thesis turns, and the retrospective audit of §3.8 tests it directly.

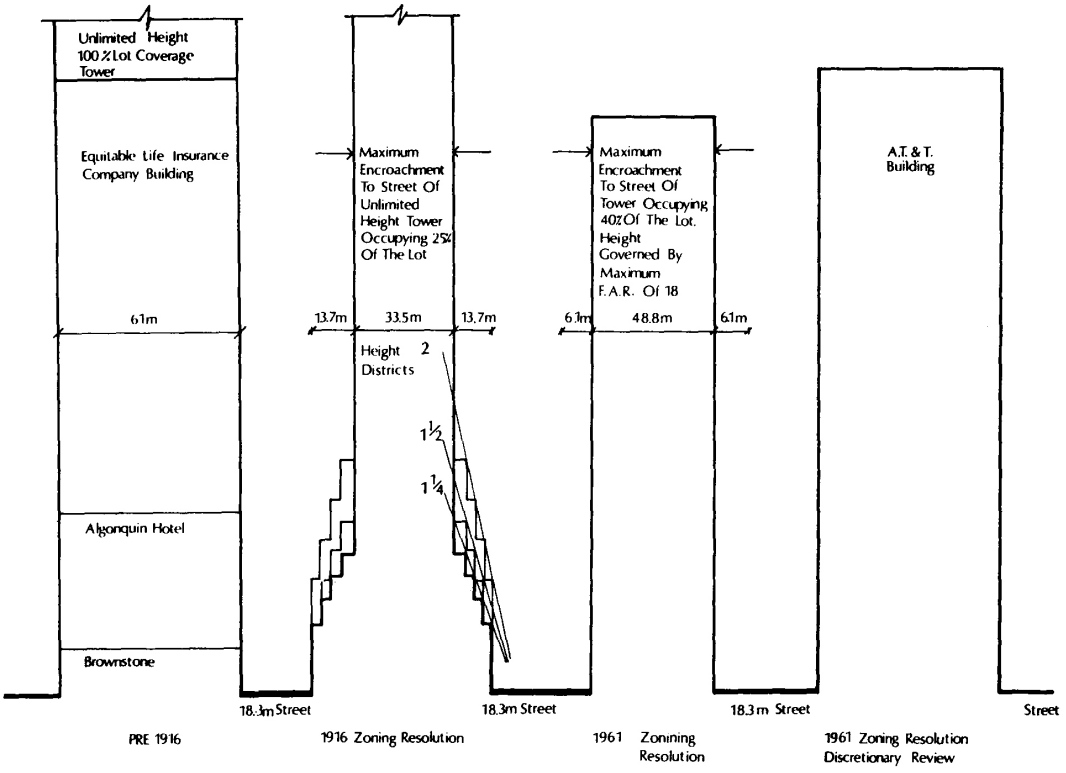


Figure 2.1 — The three regimes as envelopes drawn by the regulation's own authors: the pre-1916 unlimited height, the 1916 setback envelope of the prevalent two-times height district, and the 1961 tower-on-base under the floor-area ceiling. This is the comparison the evolutionary evaluation stages and the source of the calibrated working values recorded in Appendix A. Reproduced from Kwartler and Masters (1984, Fig. 9).

2.4 The DEC as a measuring instrument

The Chart is a curvilinear perspective projection of the sky dome as seen from the regulatory vantage point, mapped to a two-dimensional diagram whose vertical axis is distorted to compress the lower hemisphere and expand the critical region above 70° of elevation. It is divided into daylight squares, each representing 25 ft of frontage and intervals of elevation and subdivided into finer subsquares above the 70° line. A building is evaluated by projecting its corners onto the Chart, counting the squares its mass blocks, applying a profile penalty where the building encroaches beyond a drawn profile curve, and expressing the unblocked remainder as a percentage of the available daylight. The full plotting and scoring rules are specified in §§81-272 to 81-274 of the ZR.

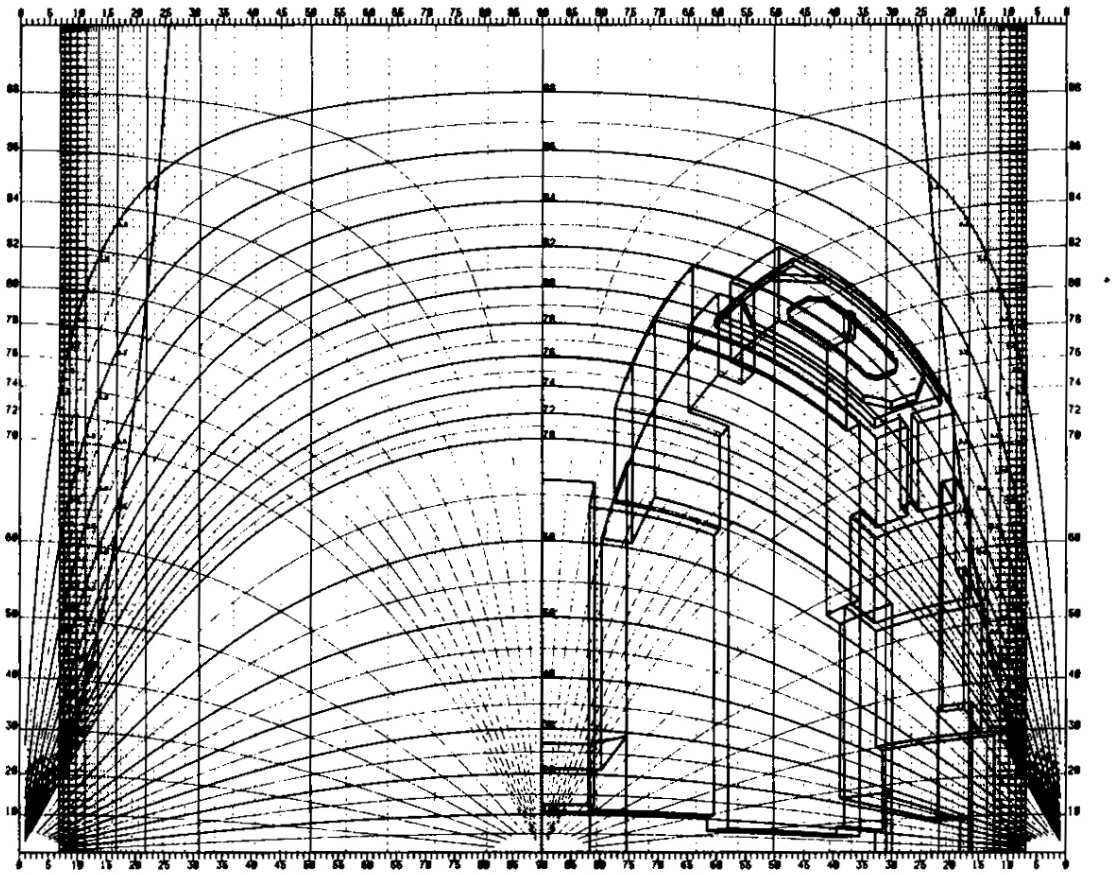


Figure 2.2 — The Daylight Evaluation Chart: a perspective projection of the sky dome mapped onto a cylinder and flattened, here with a massing plotted, following the Waldram diagram. Curved altitude bands make equal chart areas correspond to equal sky exposure but force every rectilinear building into curves that must be constructed point by point: the practical cost that motivated the 1982 reformulation. Reproduced from Fishbaut, Weingarten and Neiger (1984, Fig. 1).

The Chart's geometry made it notoriously laborious to use by hand. Fishhaut, Weingarten and Neiger (1984), architects at Skidmore, Owings and Merrill (SOM) observed that its cylindrical projection turned straight building edges into curves and made the plotting of any non-rectangular form extremely difficult, a practical constraint that, in their judgement, effectively confined the building types an architect could realistically test to simple rectangular massings, despite the regulation's stated aim of offering maximum design flexibility. Their critique is notable because it identifies a covert bias in the instrument itself: the difficulty of the evaluation, not the content of the rule, was steering design toward a narrow range of forms. They proposed a rectilinear reformulation of the Chart that reproduces its scoring while mapping straight edges as straight lines, and it is this reformulation that the present thesis implements.

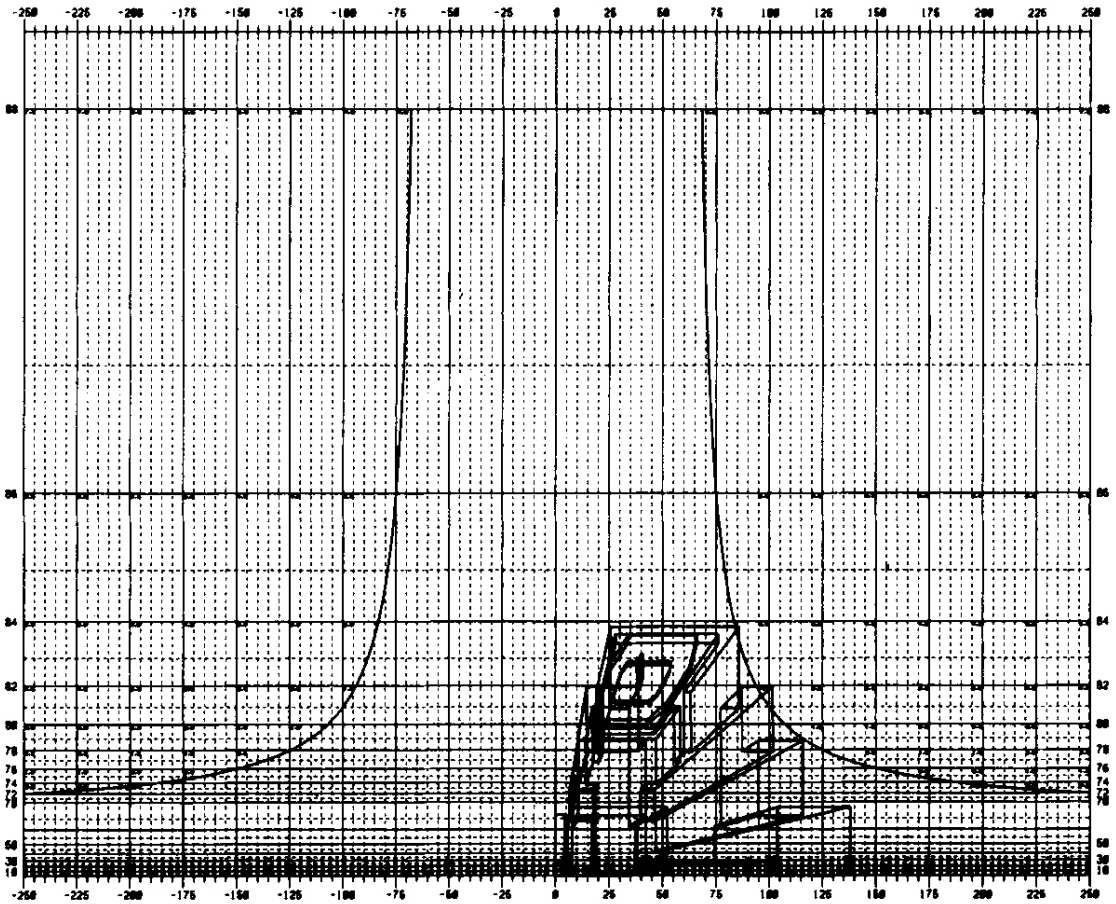


Figure 2.3 — The **SOM** *Rectilinear Chart: the rectilinear successor, spanning 250 ft either side of the vantage along the street, with the daylight profile curve and the compressed sub-70° bands visible. Straight building edges now map to straight lines, at the price of a distorted area weighting, which the subsquare values repair. Reproduced from Fishhaut, Weingarten and Neiger (1984, Fig. 9).*

2.5 The regulation since 1982

The daylight evaluation machinery remains current law. The operative sections of the ZR governing the Chart's features, plotting rules, and scoring (§§81-272 to 81-274) were last amended on 2 February 2011, in a recodification that left their mechanics substantively as adopted in 1982, and the citywide amendments of December 2024 did not touch them. The principal subsequent development is the East Midtown rezoning of 2017, which introduced modified daylight evaluation provisions, including an additional chart for qualifying Park Avenue sites; these special-subdistrict modifications are acknowledged here and excluded from the scope of the present study. In consequence, the instrument implemented and audited in the chapters that follow is not a historical artefact but the regulation in force in Midtown today.

2.6 SVF and street-level openness

The Sky View Factor (SVF) is the fraction of the upward hemisphere, seen from a point at ground level, that is open sky rather than obstruction. It runs from zero in a fully enclosed space to one in open terrain, and at intermediate values it gives a continuous, geometrically defined measure of how much sky a person standing at that point can see. It originates in urban-climate research, where Oke (1981) established the geometry of the street canyon as a determinant of nocturnal heat loss, and where it has since been used as a compact descriptor of urban texture in studies of solar access and environmental performance (Ratti, Raydan and Steemers, 2003). Its relevance here is spatial rather than thermal: it captures, in a single geometric quantity, the enclosed character of a street as experienced from the ground.

The SVF and the DEC score are geometric relatives. Both are functions of the building mass visible from a point at street level. But they differ in three respects that matter: the Chart measures from fixed vantage points at a prescribed distance, whereas the SVF may be sampled continuously along a path; the Chart weights the visible sky non-uniformly through its profile curve and 70° threshold, whereas the SVF applies no regulatory weighting of that kind, its sky weighted only by the cosine of incidence that the definition of the factor itself entails; and the Chart sweeps only 90° of azimuth, whereas the SVF integrates the full hemisphere. Two buildings may therefore earn the same DEC score while presenting a pedestrian with different amounts of visible sky. It is this potential divergence that the thesis sets out to measure. Both are, throughout, measurements of daylight rather than sunlight: neither instrument tracks the sun, and equal readings imply nothing about direct solar access.

2.7 Computational and evolutionary approaches to building form

The use of evolutionary computation to generate and evaluate building form against environmental objectives was established by Caldas and Norford (2002), who applied a genetic algorithm (GA) to facade and massing configurations under combined daylight and energy objectives, demonstrating that evolutionary search could navigate the non-convex trade-offs such problems present. Their framework operated at the scale of the individual building and treated the regulatory context as a fixed boundary condition rather than as a variable of interest.

More recent work has extended evolutionary optimisation to the scale of urban massing. Wang, Janssen and Ji (2024) present EvoMass, an evolutionary algorithm guided by diversity that generates massing from different typologies in a parametric environment, aimed at early-stage performance exploration. Liu et al. (2023) apply multi-objective evolutionary optimisation to urban block form with daylight, energy and photovoltaic objectives combined, recovering Pareto fronts that reveal the trade-offs among competing environmental goals. These studies share the same premise: evolutionary search finds good designs within a given set of rules. None of them flips the method around to treat the regulation itself as the variable and compare the design spaces different regimes produce; none measures the buildings actually built against the rules that governed them; and no implementation of the DEC tied to its published worked example was publicly available to do either. Those are the gaps the following chapters address.

3. Methodology

The investigation requires five things: a representation of the four regulatory regimes as constraints a computer can apply; a way to generate and evaluate building massing under each regime; a way to compare the resulting design spaces; an independent measure of pedestrian sky openness; and a way to score the realised building stock with the same instrument used in the experiments. This chapter sets out the method and the logic of its principal choices. The detailed specifications of the encoded constraint sets, the full Daylight Evaluation fitness algorithm, and the validation of that algorithm against the worked example in the ZR, are given in Appendices A through C and referenced here where relevant; §3.9 states what the reader needs to know of the instruments that carry the method, and the verification that they work.

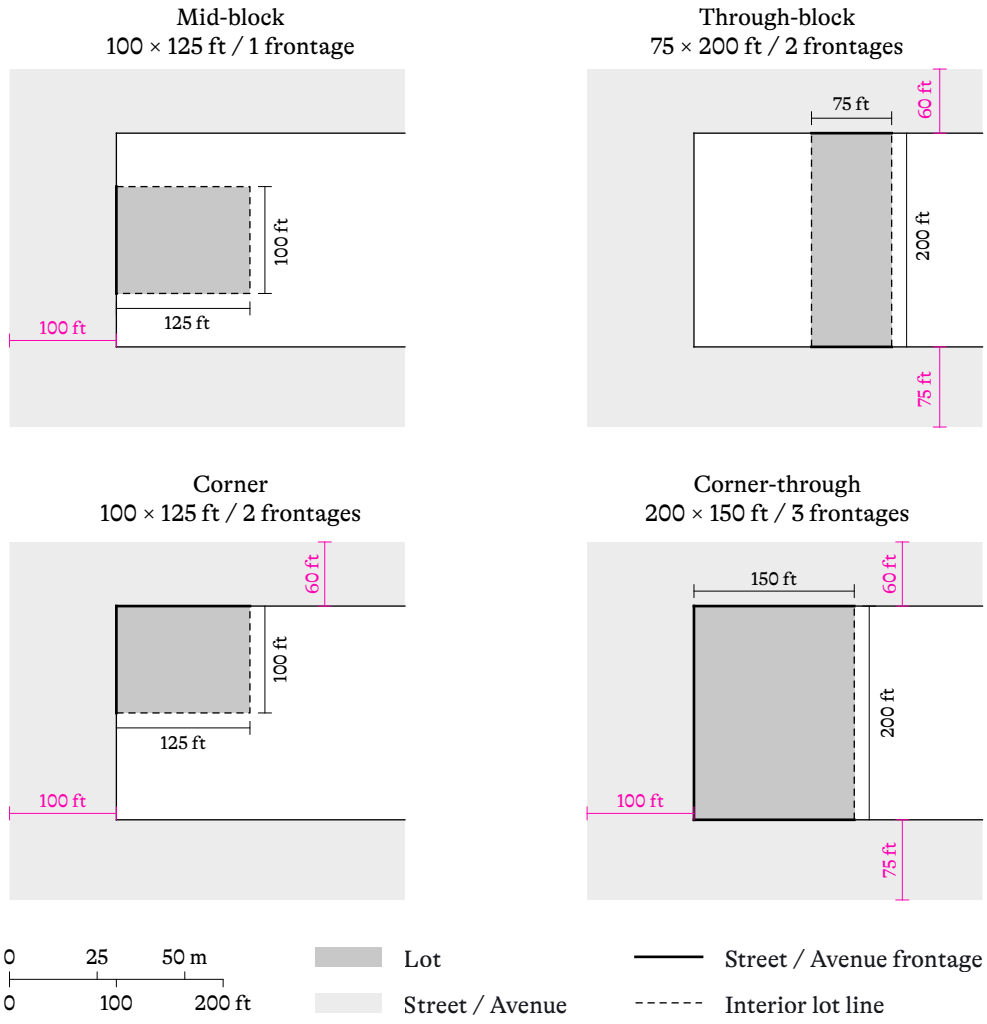


Figure 3.1 — The four canonical lot typologies of the evolutionary evaluation, in plan at a common scale, with street frontages labelled by mapped width and interior lot lines dashed: mid-block (100 × 125 ft, one avenue frontage), through-block (75 × 200 ft, two opposed street frontages), corner (100 × 125 ft, avenue and cross-street), corner-through (200 × 150 ft, three frontages). The family varies the number and opposition of frontages at comparable area, and is held fixed across all four regimes (Appendix E, LotLibrary).

3.1 Overall approach: two evaluations

The method comprises two complementary empirical evaluations that share one measuring instrument. The first evaluation is a controlled computational experiment on what each regime permits. A common set of canonical lot typologies is held fixed across the four regimes (Figure 3.1), so that any difference in the forms that emerge is attributable to the regime and not to the site. For each regime and typology, a multi-objective evolutionary algorithm evolves a population of massings against two competing objectives, daylight performance and built volume, until it approximates the Pareto front. The sixteen resulting design spaces are compared, and an independent spatial measure, the centreline SVF, is computed across the evolved populations and tested for correlation against the daylight score.

The second evaluation is a retrospective evaluation of what each regime produced. Cohorts of real Midtown buildings are scored from 3D models, sourced from NYC Department of City Planning (DCP) open data, through the same validated implementation, their distributions read against the 75% and 66% thresholds the regulation attributes to exactly this stock, and their scores compared with measured street-level openness free of any selection pressure. The first evaluation answers H2 and contributes to H3; the second answers H1 and provides the test of H3 (§3.8).

Two objectives are used rather than three. Daylight performance and built volume stand in direct tension: taller, bulkier massing blocks more sky, while higher scores demand the setback and slenderness that cost floor area, and that tension is what makes the optimisation a meaningful search rather than a single best answer. Volume is adopted as the proxy for development capacity: for massing studies it is monotonically related to floor area and requires no contestable model of rents or construction cost. Two objectives also yield a legible two-dimensional front. The SVF is excluded from the objectives by design and reserved as an independent measure (§3.5).

3.2 Representing the regimes as constraints

Each regime is encoded as a constraint module determining which forms are admissible: no limit for the pre-1916 baseline; street-wall height, sky exposure plane and tower coverage for 1916; a FAR, coverage limits and the tower-on-base envelope for 1961; and the DEC itself, as a performance constraint, for 1982. Massing is generated on a footprint distinct from the lot boundary, so that the street wall of the first regime and the plaza of the second emerge from the optimisation rather than being prescribed to it.

Because each regime varies several regulatory dimensions at once, differences between evolved populations cannot be attributed to any single rule; the protocol of §3.6 addresses this confounding.

Each module is realised as a compiled C# component and reports a graded compliance score rather than a binary verdict; how legality enters the fitness (enforced by construction under the prescriptive regimes, a graded threshold penalty under the 1982) is the subject of §3.7. Appendices A.6 and A.7 give each module's rules and the treatment of legality; A.3 records the parameter values and the modelling of the plaza bonus.

3.3 The Daylight Evaluation fitness function

The daylight objective is a faithful implementation of the DEC as specified in §§81-272 through 81-274 of the ZR. Rather than reproduce the Chart's cylindrical projection, it adopts the rectilinear reformulation of Fishhaut, Weingarten and Neiger (1984), equivalent to the DEC for scoring purposes and far more tractable to compute, preserving the number and relationship of the daylight squares and admitting an efficient implementation by rasterisation. Its validation against the worked example given in the ZR, the same example published by Kwartler and Masters (1984), is reported in §3.9 and Appendix C.

The instrument itself is a purpose-built pair of components: DECProjection and DECScorer. Two vantage points are constructed on each street centreline per the plotting rules of §81-273; the projection maps every massing vertex onto the chart plane by the rectilinear mapping, reducing blockage testing to polygon–rectangle intersection; the scorer rasterises the chart surface into subsquares, charging a subsquare wherever the silhouette covers any part of it, the strict any-portion rule of §81-274(a)(3). The profile penalty is triggered by the profile curve, which the published text does not give in closed form; it is encoded as a width-specific digitised coordinate table with linear interpolation, traced from the NYC DCP published diagrams reproduced in Appendix C.4. Vantage scores are averaged per street, street scores combined by frontage-weighted average into the overall score, and the per-street minimum carried separately for the 66% test. The scorer is a pure function from massing, lot and street widths to a score, which is what allows the identical compiled component to serve both evaluations. Appendices C.5 and C.6 specify the projection, rasterisation, credit rules and aggregation in full.

3.4 The evolutionary algorithm

The optimisation uses NSGA-II (Deb et al., 2002), which ranks individuals by Pareto dominance and preserves diversity along the front by a crowding-distance measure, returning not a single optimum but the whole set of trade-off solutions. That property suits a study comparing entire design spaces rather than seeking any one best building. Regimes are compared by the hypervolume indicator (Guerreiro, Fonseca and Paquete, 2021), which expresses in a single number how much of the achievable daylight–volume space each front dominates against a fixed reference point. The genotype and the treatment of legality are specified in §3.7; population, operator and convergence settings are reported with the experimental results.

The optimisation runs inside the Rhino–Grasshopper parametric environment, driven by the Wallacei implementation of NSGA-II (Makki, Showkatbakhsh and Song, 2024) over the pipeline of single-responsibility components inventoried at §3.9: pre-processor, phenotype builder, regime modules in parallel, scorer closing the loop. All geometry is held in metres, with street widths and chart coordinates in feet, the unit regime of the regulation itself.

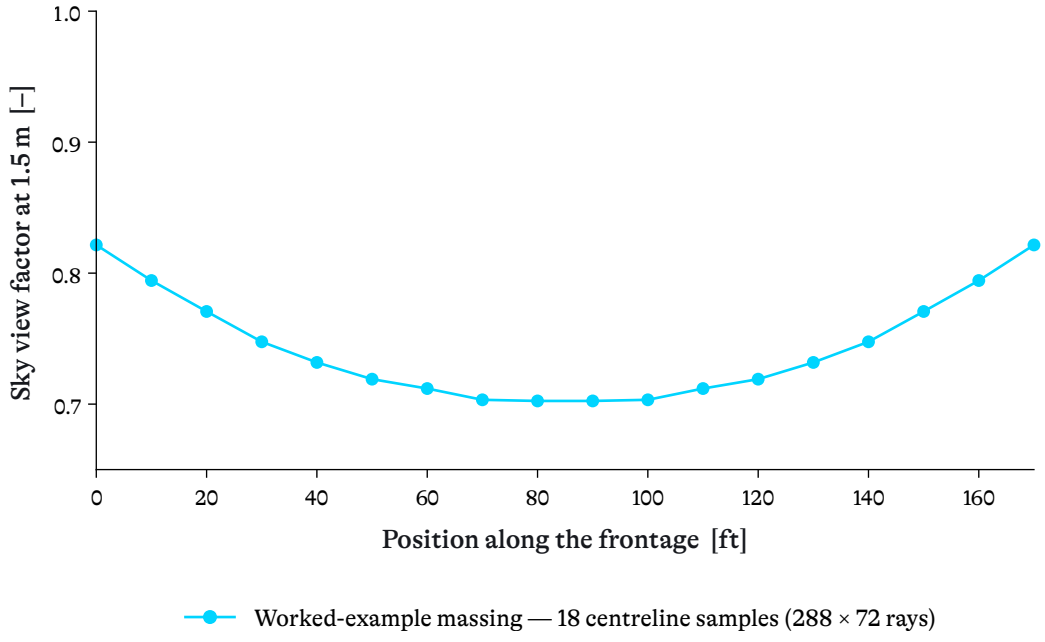


Figure 3.2 — Along-frontage sky-view profile of the worked-example massing (§3.9), sampled at pedestrian eye height (288 × 72 rays): eighteen samples, symmetric to six decimal places, minimum 0.7024 at the centre. The Chart's two fixed vantage points structurally cannot record this variation; the profile is what the SVF instrument adds.

3.5 The SVF as an independent measure

The centreline SVF is the mean sky view factor along the street centreline at a pedestrian eye height of 1.5 metres, sampled at regular intervals across each lot's frontage.

It is measured by a purpose-built sampler, CenterlineSVF, which discretises the upward hemisphere at each sample on an altitude-by-azimuth ray grid and reports the unobstructed fraction (Figure 3.2); the sampling specification and resolution policy are given in Appendix C.2. The rays test the evolved or modelled building alone, with no surrounding context massing; this is the convention the DEC itself adopts, since the Chart likewise plots only the building under evaluation, and it ensures that any divergence between the two measures is attributable to how each instrument samples the sky, not to what each is allowed to see. The SVF is held apart from the optimisation by design: it is never an objective, and no building is ever selected for it. Its role is diagnostic, measured after the runs across the resulting populations and, in the retrospective evaluation, across cohorts where no selection pressure operates at all. The geometric reasons the two may diverge are set out in §2.6.

3.6 Experimental protocol and the confounding problem

The score-to-SVF correlation is at risk of circularity in the 1982 regime, where the population has been bred for daylight performance: a high SVF there may reflect selection for a correlated quantity rather than a genuine coupling between the two measures. The protocol measures the correlation across all four regimes and foregrounds the three in which daylight was not the selection objective, the pre-1916, 1916 and 1961 regimes, as the uncontaminated test within the evolutionary evaluation. The retrospective evaluation supplies a second, independent resolution: real buildings were designed under economic and architectural pressures, not bred by the experiment, so the correlation measured across the cohorts is free of the selection artefact by construction.

The correlation is measured on a drawn sample rather than the whole population, constructed per cell to cover the daylight axis evenly: distinct genomes are taken once, the individuals on the Pareto front are set aside as a separate on-front sample, and the remaining off-front individuals are stratified by DEC score into fifteen equal-width bins with up to eighteen drawn from each. The off-front set is the primary evolutionary test; the on-front sample is reported alongside it. Appendix C.8 specifies the protocol, the stratification and the alignment checks.

3.7 Genotype, decoding, and the treatment of legality

A building massing is encoded as a fixed-length vector of real-valued genes, decoded as a cumulative sequence of extrusion-and-setback operations. The genome opens with leading footprint-inset genes, one per street frontage class, setting how far the ground storey stands back from the lot line. These are followed by one group per segment: a rise gene, the height of the next extrusion step, and one inset gene per tagged setback frontage class, the depth the mass steps inward before the next rise. Setbacks are true inward polygon offsets of the active footprint with a general mitre-and-reclosure rule so that the same decoding holds across all lot typologies. Because insets may decode to zero, the number of setbacks an individual exhibits is an emergent property of its genes rather than a stored parameter, which keeps the genome fixed-length and the crossover operator meaningful.

The massing rises from a building footprint contained within the lot; the lot polyline remains the reference for every regulatory measurement (street walls, coverage ratios, the vantage construction of the Chart), so that coverage ratio and plaza area fall out of the representation as free descriptive statistics. Up to eight segments are decoded, each rise gene mapping to a bounded storey-height increment and each inset gene to a bounded inward offset applied per tagged frontage class, so that a corner lot may set back differently from its avenue and its side street.

Regime legality enters the fitness differently for prescriptive and performance regimes, and the asymmetry is a deliberate design decision. The prescriptive regimes (1916 and 1961) define legality geometrically and the pipeline enforces it by construction: the decoder clips every extrusion to the regime envelope, so a non-compliant phenotype never enters the population, and the legality factor those modules report is one throughout, an audit channel rather than a penalty, verified across every individual sampled in §4.2. The 1982 regime defines legality as a score threshold and applies a soft penalty proportional to the shortfall against the 75 and 66% thresholds; its effect on H2 is audited in §4.2. That its feasible region is bounded by a performance surface while the earlier regions are bounded by geometric envelopes is not a modelling convenience but the historical substitution under study, namely a score in place of an envelope, which is what the 1982 Resolution enacted; the comparison of design spaces in H2 is therefore a comparison across that substitution. The objective framing, the negation convention and the placement of the penalty on the volume objective are set out in Appendix A.7.

3.8 Retrospective evaluation of the built record

The second evaluation scores cohorts of real Midtown buildings, five constructed under each of the 1916, 1961 and 1982 regimes, from 3D massing models (Figure 3.3), using the same validated implementation as the experiments, with each building's lot geometry and street widths reconstructed from city records. For the 1961 regime these five constitute the entire as-of-right Midtown population the regime produced, a sparseness that is itself a finding rather than a sampling limitation. For the 1916 and 1961 cohorts the exercise is a calibration audit: the regulation states that its thresholds were derived from the measured performance of exactly this stock. For the 1982 cohort it is an evaluation rather than an audit, since for buildings that did not elect the evaluation route the Chart serves as a uniform measuring instrument, not a legality test, and a score below 75% is not a finding of non-compliance.

Cohort selection applies four gates before any model enters the study: as-of-right construction under its regime, excluding the discretionary-approval cases whose departures from the envelope motivated the 1982 reform; no single street frontage exceeding 250 ft, the geometric limit beyond which §81-275 applies; a lot geometry admitting the regulatory vantage construction; and records documented well enough to model without conjecture. Appendix B.1 records the screen, and every exclusion with it.

The audit's acceptance criterion accounts for the measurement conventions of the instrument. The implementation reads conservatively relative to 1984 hand drafting: on the regulation's own worked example it returns 72.89% against the published 76.75%, a residual that is fully decomposed in Appendix C and biased in a single, known direction. Cohort scores are therefore reported raw and read against the worked example as an anchor: a stock genuinely calibrated to the published 75% would read approximately 71 under this implementation (75 less the 3.86-point conservative residual); the acceptance band is set at [68, 76], its floor three points below the anchored expectation and its ceiling one point above the published threshold, so that a cohort consistent with either reading of the calibration falls inside it. Cohort means falling within that band are consistent with the 75% claim, while means materially below it falsify H1. No bias correction is applied, because the residual is convention-driven rather than a constant offset; it is reported and bounded instead. The analysis itself is deliberately compact: one score distribution per regime read against the 75 and 66% lines, and one correlation analysis of DEC score against measured centreline SVF across the pooled cohorts, the latter serving H3 as the selection-free replicate of the evolutionary evaluation's correlation.

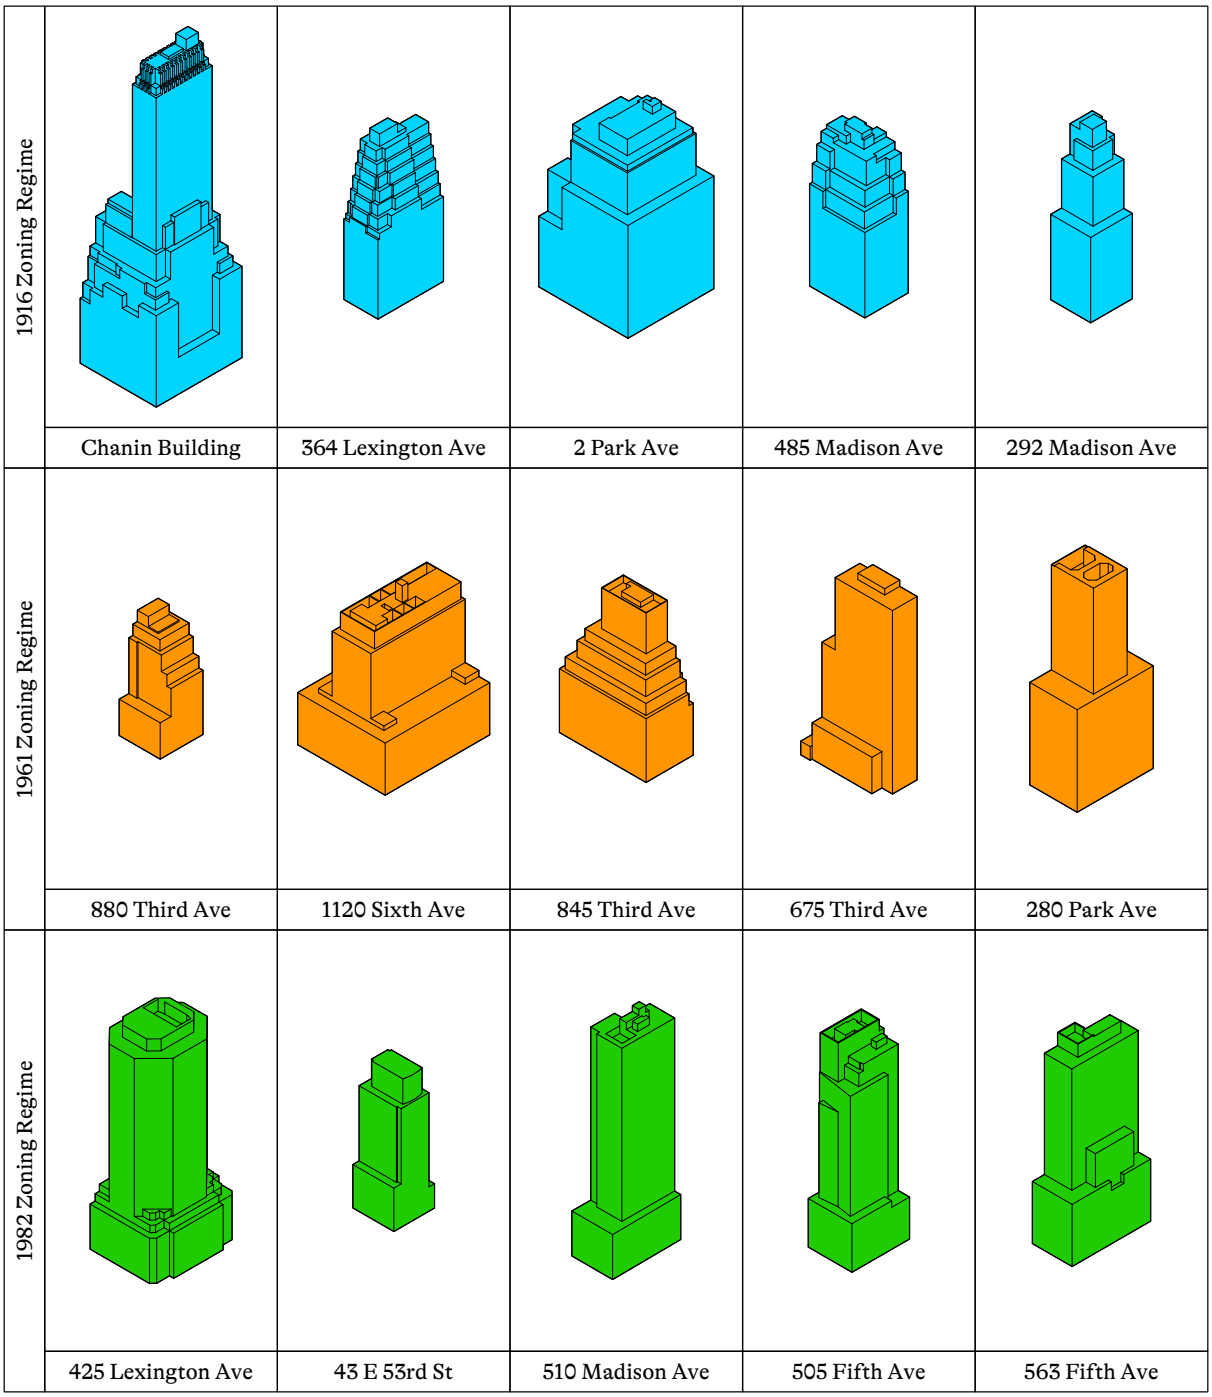


Figure 3.3 — The retrospective cohort as modelled: the fifteen as-of-right buildings, five per regime, in axonometric projection at a common scale. The regime typologies are legible in the massing: the 1916 setback profiles, the 1961 towers on base, the 1982 street-wall hybrids. Identifications and tax lots at Table B.1.

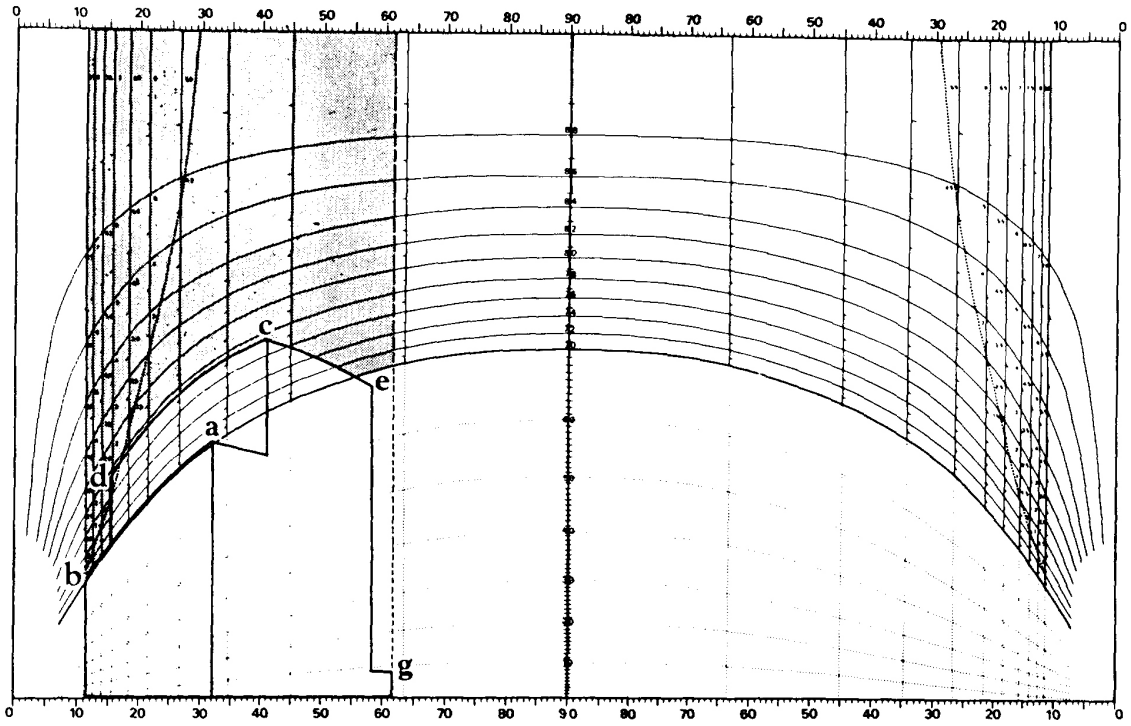
The retrospective pipeline reuses the scorer without the evolutionary harness: each building enters as a massing mesh with its tagged lot geometry, and the same model is sampled for the centreline SVF. Appendix B records the screen and the cohort: the 1916 five drawn from ninety-five regular-lot candidates, the 1961 five as noted above, the 1982 five as-of-right under the regime, their elected compliance routes unrecoverable from public records and discussed at B.8.

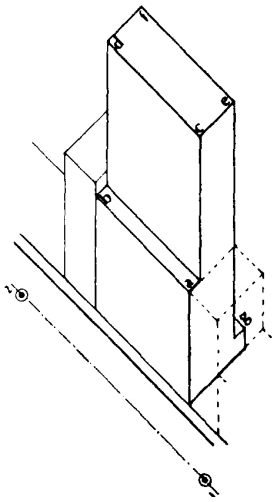
3.9 The instruments

The method is carried by purpose-built C# components for Rhino–Grasshopper: a lot pre-processor; the phenotype builder of §3.7; the four regime constraint modules of §3.2; the DEC projection and scorer of §3.3; the centreline sky-view sampler of §3.5; and an HV calculator implementing the indicator of §3.4. Four facts govern the reading of every number here. Both evaluations share the identical compiled scorer, so a percentage point means the same thing in the experiments and the audit. The scorer's counting convention, the strict any-portion rule of §81-274(a)(3), is declared because, as the validation below shows, no DEC score is comparable with another unless both declare the rule by which a subsquare was charged. The sky-view sampler is never an objective and is measured only after selection has finished. And each instrument was verified before any experiment was run, against the regulation's worked example and against closed-form expectations on synthetic geometry; every implementation detail is confined to Appendices A, C and E.

The scorer was validated against the worked example published in §§81-273 and 81-274 of the ZR, the same example scored by hand in Kwartler and Masters (1984): a three-slab massing of street wall, tower and low rear block on a 170-by-100-foot lot fronting a 100-foot street (Figures 3.4 and 3.5). A rig component constructs it programmatically, and the symmetry of the massing supplies an internal check: both vantage points must return identical scores.

The validation decomposes the score into its regulatory components (Table 3.1). The available-daylight term matches to within a tenth of a square and the below-seventy credit exactly; the blockage term reads 23.7 squares against the published 20.5, and the profile penalty 0.70 against 0.4, for 72.89% against the published 76.75. The two components carrying the gap, blockage and profile, are both conservative, and the aggregate never reads higher than the published practice, which is what fixes the acceptance band of §3.8. The reconciliation is in Appendix C.7.





The example building should be scored as follows:

AVAILABLE DAYLIGHT SQUARES	89.9
SQUARES BLOCKED TIMES VALUE	
Squares above 70° blocked	−20.5
Squares below 70° open	+0.0
Profile Encroachment	
0.1 × 1.5 / 0.25 × 1 =	−0.4
Total blockage	−20.9

DAYLIGHT SCORE

$$\frac{89.9 - 20.9}{89.9} = 0.7675 = 76.75\%$$

Figure 3.4 — The worked example: Kwartler and Masters (1984, Fig. 10), the hand-drafted plot on the curvilinear chart with its printed score sheet (available 89.9, blockage −20.5, profile −0.4, overall 76.75%). Figure 3.5 plots the identical massing (a 170 ft frontage on a 100 ft street carrying a 140 ft street wall, 320 ft tower, and 23 ft rear slab) through this study's rectilinear instrument; the two projections differ by construction, and Table 3.1 reconciles the arithmetic term by term.

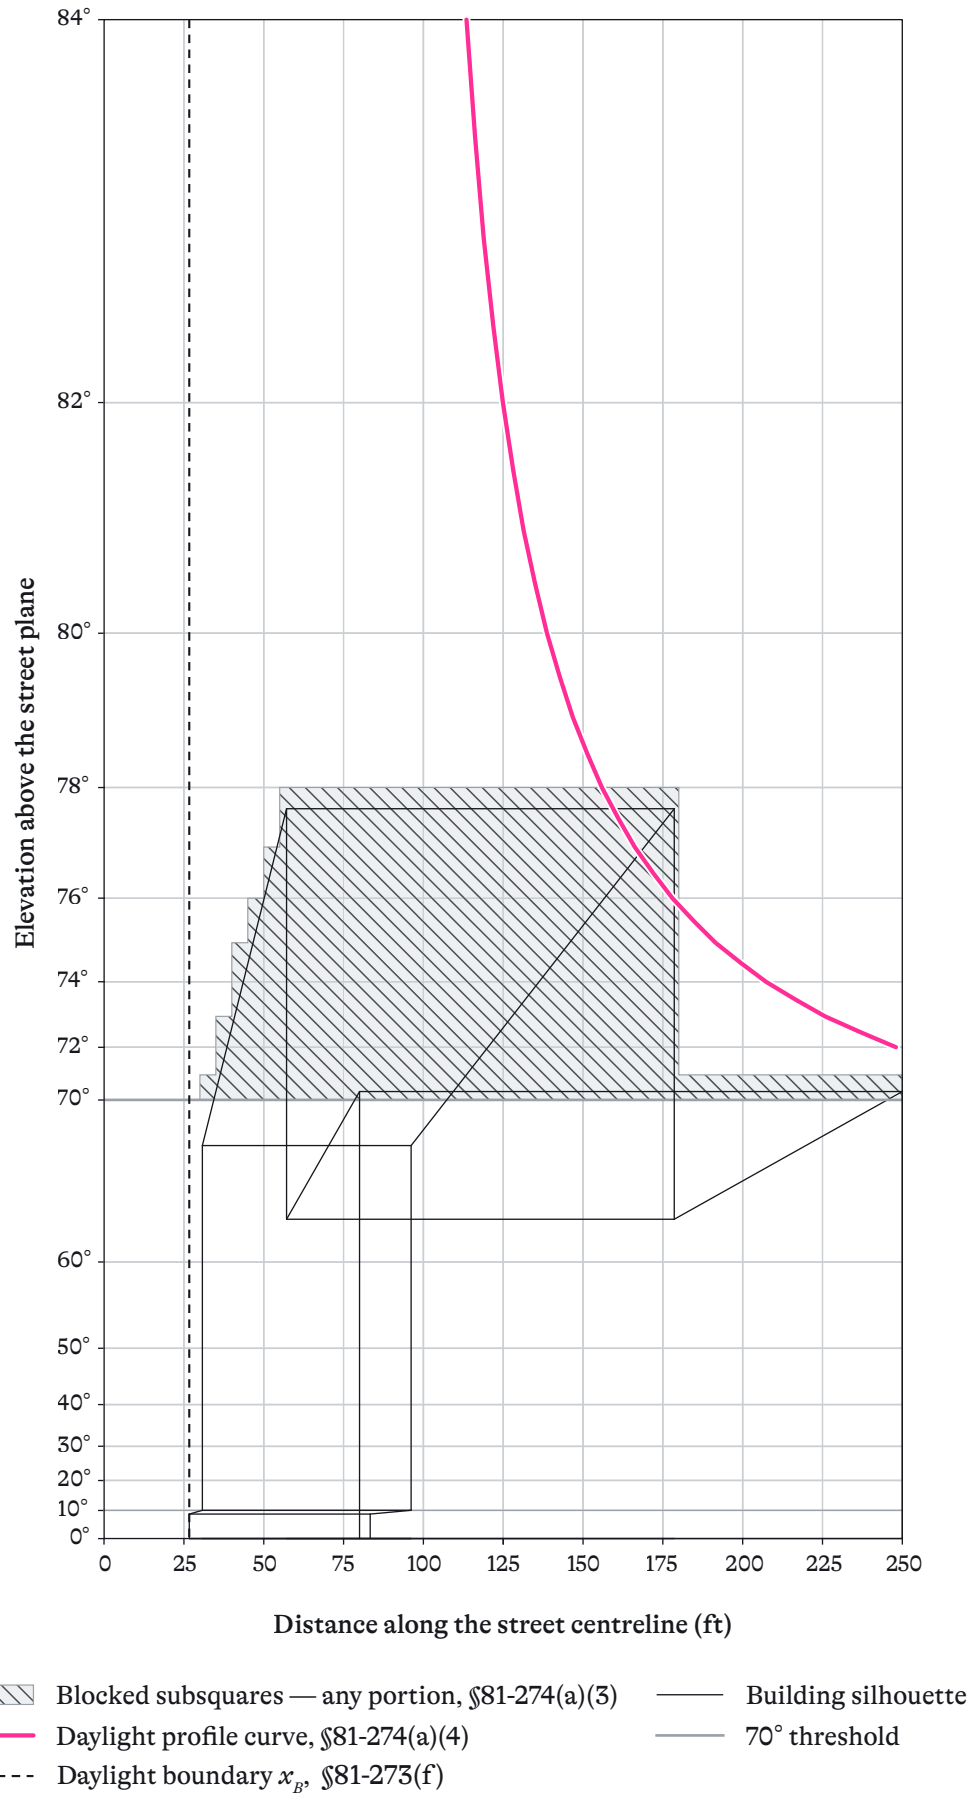


Figure 3.5 — The worked-example massing, scored by this study's instrument. The vantage projects at $x = 0$; squares are counted from the daylight boundary ($x_B = 26.7$ ft) to the far lot line at 250 ft.

The residual is not merely arithmetic. Under the canonical rule the regulation's own worked example scores 72.89% and so fails the 75% threshold it was published to demonstrate, while passing the 66% per-street test. Two readings of one sentence of §81-274(a)(3), charging a subsquare when the silhouette covers any part of it or only its centre, separate a pass from a failure on the very building the ZR uses to teach the method. Neither is arbitrary: any-portion is the text's own rule, and that of the algorithm SOM published to implement it; centre-point reproduces the tolerance a hand-drafter exercised on grazing rows. The chart is therefore not fully specified by its published arithmetic: a daylight evaluation is a score plus a counting convention, and the convention is nowhere written down. This is a finding about the regulation, not the software; its consequences are bounded in C.9.

Verification extends beyond the worked example through closed-form unit tests whose rigs generate geometry and compute the expected value analytically; they never re-implement the instruments they test. For the scorer: a negligible mass returns 120.6%, above one hundred by design, the credit adding to available daylight; a mirror-symmetric slab returns identical vantage decompositions; a 120-foot wall exercises the otherwise-dead credit loop (+4.2 of a possible +16.8); a pole at the lot line confirms the rectilinear identities to 4×10^{-6} ft and the 50/51 depth scaling to eight digits; and a three-frontage corner-through rig confirms the six-vantage, length-weighted aggregation of §81-274(h) in closed form. The symmetry test failed on its first geometry, and the failure was a finding: single-precision mesh storage at a chart gridline turns last-bit arithmetic into whole-column differences under the closed-interval counting rule, a measure-zero tie, documented at C.5; off the gridline the test passes exactly. The HV calculator carries its own analytic self-test (Appendix C.3).

Scoring component	Published (ZR / K&M, 1984)	Implementation
Available daylight	89.9	90.0
Blockage above 70°	−20.5	−23.7
Credit below 70°	0.0	0.0
Profile penalty	−0.4	−0.7
Overall score	76.75%	72.89%

Table 3.1 — DEC scorer validated against the ZR worked example (Kwartler and Masters, 1984). The aggregate residual is conservative: the implementation never reads higher than the published practice.

For the sky-view sampler the programme was decisive twice over. An empty scene returns 1.000000 exactly, the self-normalisation that licenses second-decimal differences. An infinite canyon at two aspect ratios, H/W ≈ 1.7 and 0.2, expected 0.30 and 0.96, matches Oke's (1981) eye-corrected closed form within 5×10^{-3}, a test that first falsified the rig, whose expectation had carried the circular-well formula in place of the canyon's. A symmetric massing mirrors to six decimal places at every resolution. And the convergence test failed its pre-registered criterion: doubling the 72-by-18 grid moved individual samples by up to 0.013 and building-level means by 0.002. The instrument was rebuilt rather than excused: the retrospective evaluation is measured at 288-by-72 rays, verified against the analytic anchors to within 5×10^{-4} at the mean, while the evolutionary evaluation retains 72-by-18 after re-measurement of the most exposed cell moved its per-individual means by at most 0.0023, its rank correlation by 0.004, and its legal-minimum sky by 0.0005 (C.2; §5.3).

4. Results

4.1 H1 — calibration: cohort DEC scores

The fifteen buildings were scored through the implementation of §3.9, each frontage at its documented street width and the vantage the ZR prescribes (Figure 4.1). The five 1916 buildings return a mean DEC score of 56.9% (44.9 to 69.2), and only 292 Madison Avenue reaches the acceptance band, at its floor. The five 1961 buildings return a mean of 87.5% (76.0 to 97.7), every one clearing both thresholds. The 1982 buildings, designed against the chart, cluster at 79.5% (78.1 to 82.0, sd 1.5); 505 and 563 Fifth Avenue lie in the Fifth Avenue Subdistrict (FAS) and are retained under that flag (B.8).

H1 is not supported. The regulation's claim that the calibrating as-of-right stock scores approximately 75% fails on both prescriptive regimes, in opposite directions. The 1916 towers sit 18.1 points below the nominal line and some 14 below the anchored expectation of §3.8, falsifying the hypothesis by its own criterion; the 1961 towers sit 12.5 and 16 points above, a departure the criterion does not name but which contradicts the claim no less. That the 1982 cohort clears the band from above, its five scores spanning under four points, is the control.

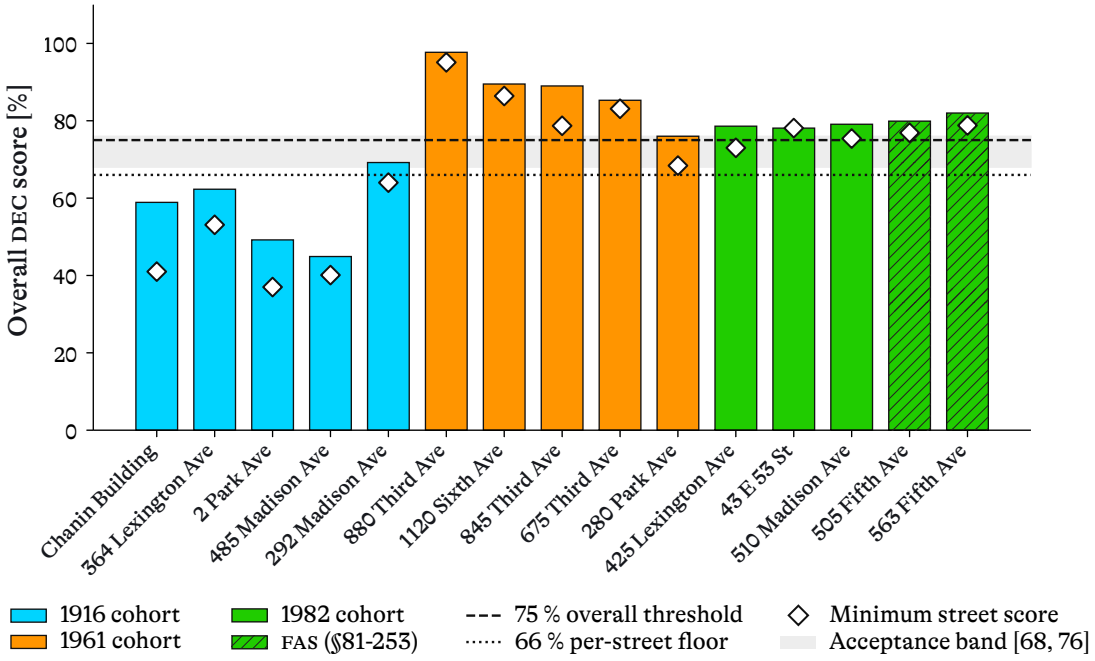


Figure 4.1 — DEC scores of the retrospective cohort, by regime: the overall score for each building. Diamonds give each building's minimum street score, against which the 66% per-street test is read; the shaded band is the acceptance band of §3.8, [68, 76]. The two hatched 1982 buildings lie in the Fifth Avenue Subdistrict and were not held to the unmodified chart (B.8).

One qualification bounds the verdict, and it bears on both limbs. All five 1916 buildings in the cohort are corner lots. The second limb, that the worst-case interior-lot building under the 1916 rules scores approximately 66%, is therefore untested here, the very case the floor was set by being absent from the cohort. The first limb is bounded by the same fact: §4.4 traces the 1916 shortfall to cross-street frontages that collapse where a sheer wall on a 60 ft street fills the vantage, a property of corner geometry on narrow cross-streets as much as of the regime. The falsification is accordingly scoped to the as-of-right corner-lot stock the cohort contains. The vantage is not a second qualification: the three 1916 frontages on 75 ft streets are scored at the 40 ft centreline the ZR fixes, not the geometric half-width, a correction worth 1.0, 1.6 and 3.1 points, enough to carry 292 Madison Avenue across the band floor. The per-street test sharpens the reading: all five 1916 buildings fail the 66% minimum as well as the overall threshold, four on a 60 ft cross-street, 292 Madison Avenue on Madison Avenue itself.

4.2 H2 — envelope: HV across regimes

The evolutionary evaluation produced sixteen Pareto fronts, one for each pairing of the four regimes with the four lot typologies, every front the outcome of an identical optimisation: one hundred generations of one hundred individuals from a single seed, at Wallacei's default operator settings, against the negated daylight and volume objectives of §3.7. Table 4.1 reports the HV each front dominates as a fraction of the unconstrained pre-1916 front on the same lot, so that every entry is the share of the unconstrained design space its regime preserves (Figures 4.2–4.4).

The ordering is the same in every typology column: the unregulated baseline admits the most form, the 1982 chart recovers most of it, and the two prescriptive regimes fall well behind, the 1961 furthest. H2 holds: at matched daylight the 1982 regime admits a materially wider space of viable form than either prescriptive predecessor, within the sixteen-cell harness of this study.

	Mid-block	Through-block	Corner	Corner-through
Pre-1916	1.0000	1.0000	1.0000	1.0000
1982	0.8630	0.8483	0.7502	0.7421
1916	0.5905	0.4747	0.5521	0.5582
1961	0.3510	0.3695	0.4458	0.4235

Table 4.1 — Pareto-front HV by regime and lot typology, each given as a fraction of the unconstrained pre-1916 front on the same lot. The baseline is unity by construction; rows are ordered by the typology-invariant rank.

The margin must be read with the asymmetry of §3.7 in view. The prescriptive regimes bound their feasible regions with geometric envelopes; the 1982 regime bounds its region with a performance surface graded around the threshold. Because the daylight score is both first objective and legality test, the 1982 front can extend to forms an envelope would have excluded outright, so part of its advantage in Table 4.1 may follow from the kind of constraint, not only its severity. The evolved record shows how much: under the 1916 and 1961 envelopes every sampled individual carries a legality factor of exactly one; the 1982 fronts are free to trade across the legality line, and barely do. Three of the four hold no sub-threshold form (left extents 75.46 on the corner, 76.06 on the mid-block, 75.15 on the through-block); the corner-through holds one, at 74.33. Removing it moves that cell by 0.117% of its own HV, its ratio falling from 0.7421 to 0.7412, and leaves the other three unchanged. The soft constraint shapes the edge of the 1982 design space without generating its width; the margin survives intact.

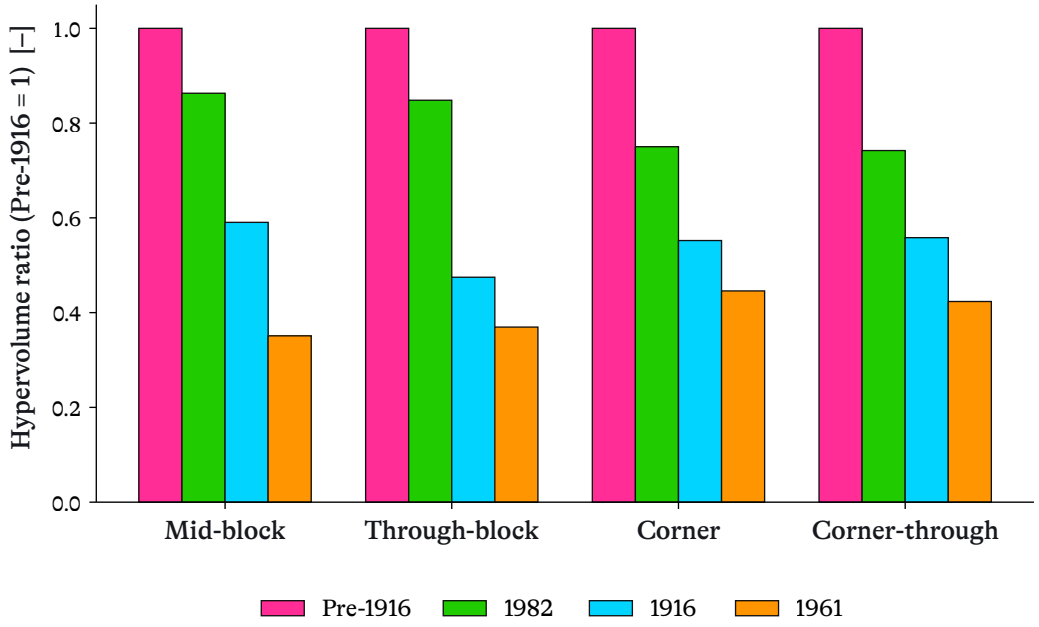


Figure 4.2 — HV ratios by regime and typology, as grouped bars: the ratios of Table 4.1, one cluster per typology; the rank pre-1916 > 1982 > 1916 > 1961 recurs in every cluster.

A further pattern in Table 4.1 speaks to the regimes' character rather than their rank (Figures 4.5 and 4.6), dividing the three regulated regimes three ways. The 1916 setback answers to what the streets are: most permissive on the mid-block (0.5905) and the avenue-bearing corner-through (0.5582), where a 100 ft frontage lets a tall street wall stand before the sky-exposure plane bites, and most restricted on the through-block (0.4747), whose two opposed planes meet in a steep tent. The 1961 envelope answers to how many street lines the tower rule can draw on, peaking on the corner (0.4458) and bottoming on the mid-block (0.3510), nearly the reverse ordering. The 1982 chart moves against the regime it replaced, its ratio falling as frontages multiply and couple, from 0.8630 on the mid-block to 0.7421 on the corner-through, where 1916 posts its near-maximum; its absolute spread, 0.121, the widest of the three (1916, 0.116; 1961, 0.095), though the smallest against its own mean. Prescriptive severity tracks the width and opposition of frontages, performance severity their number and adjacency: the substitution redistributed, rather than flattened, the regulation's geographic incidence.

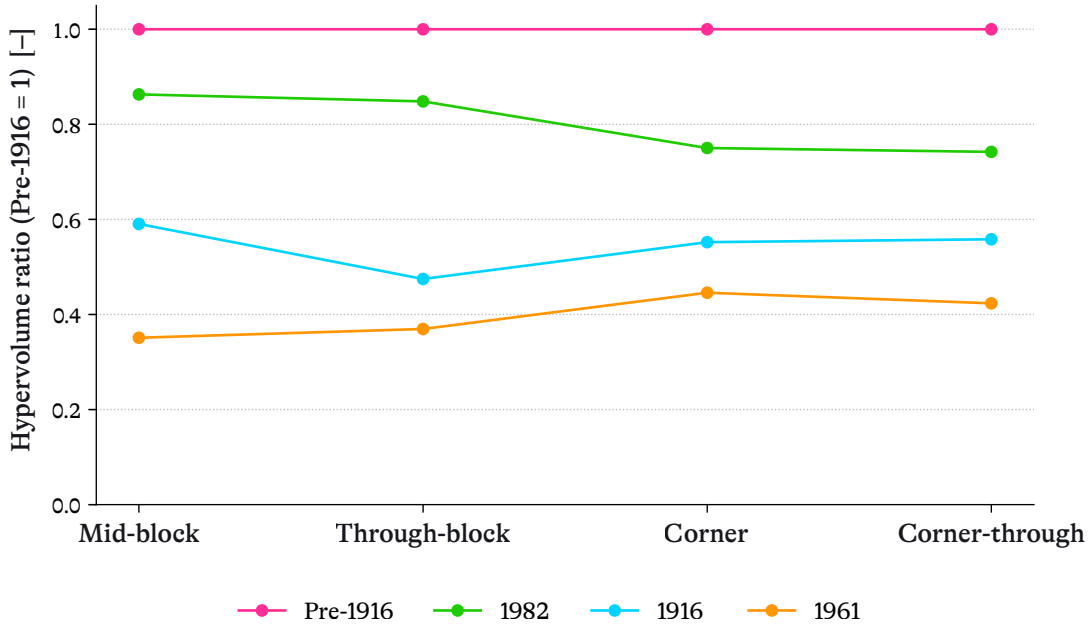


Figure 4.3 — HV ratio across the four typologies, one line per regime; the two prescriptive regimes lift on the avenue-bearing lots while the 1982 regime declines as frontages multiply and couple.

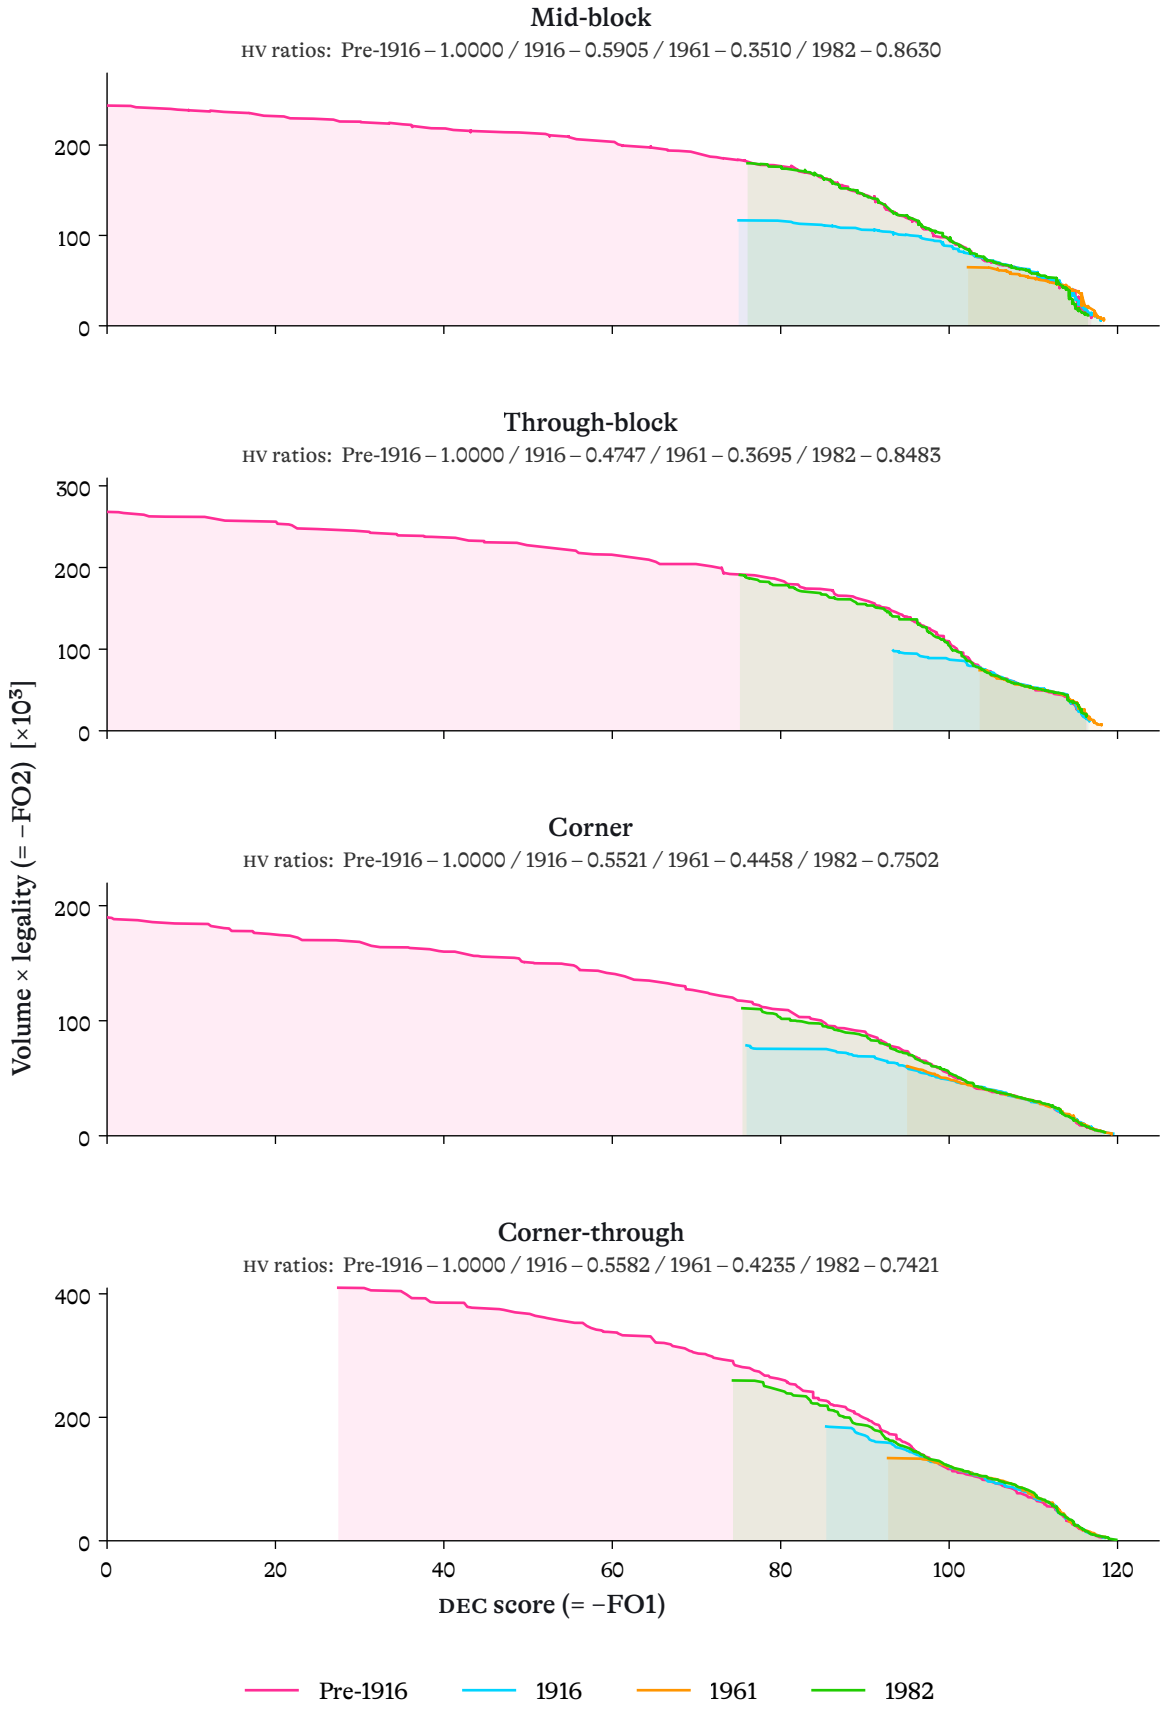


Figure 4.4 — Pareto envelopes by regime for each of the four typologies; shaded areas approximate the dominated HV, with the ratio to the pre-1916 baseline per panel (Appendix D.1).

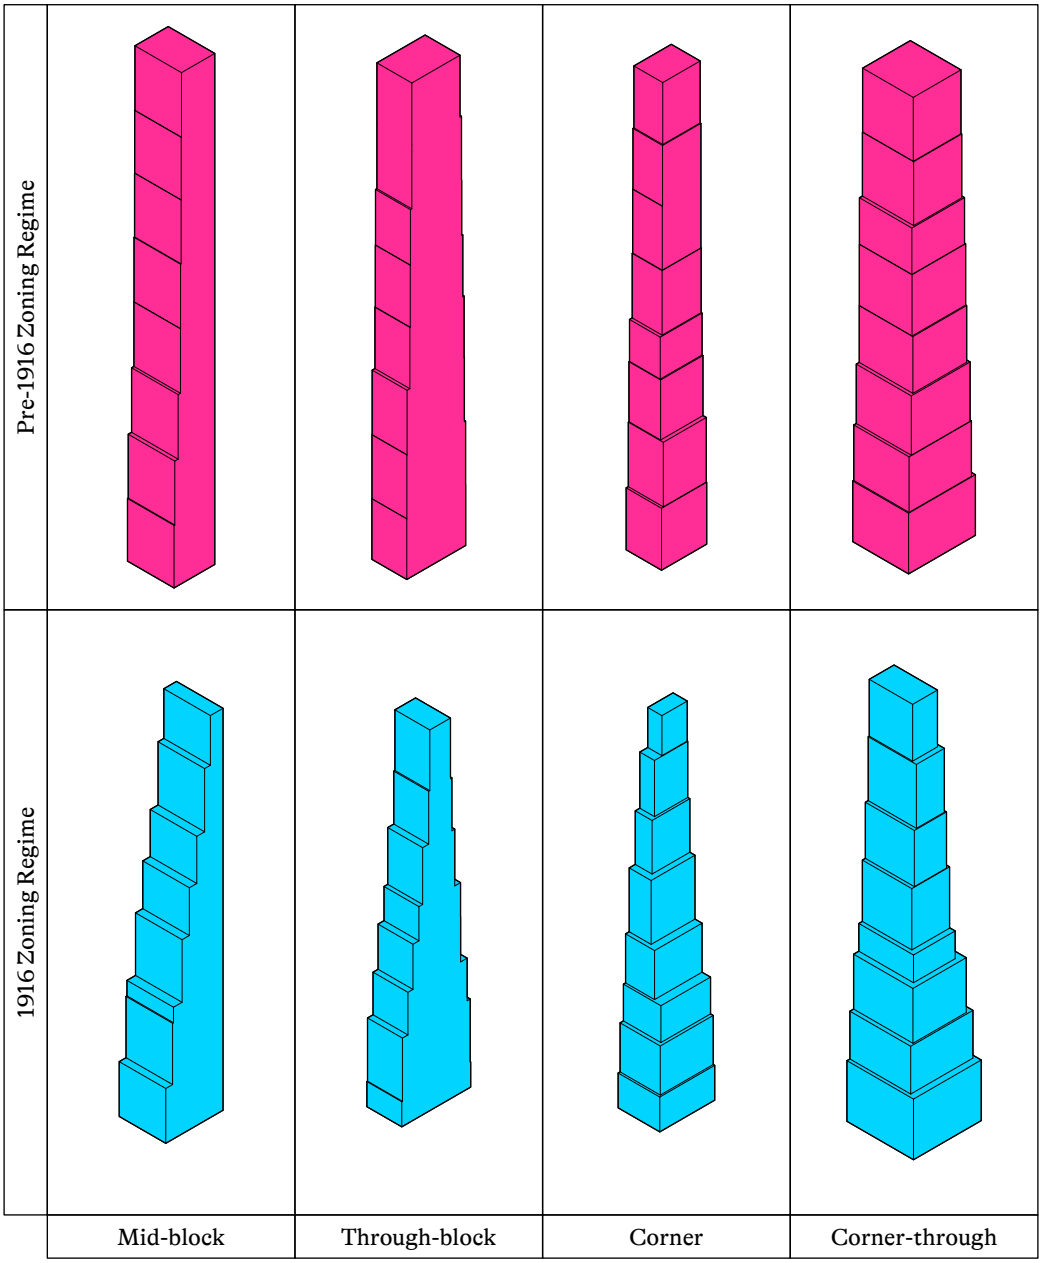


Figure 4.5 — One representative individual per cell, I: the pre-1916 baseline and the 1916 regime across the four lot typologies. Each massing is the equal-weights Chebyshev selection from its cell's certified Pareto front, drawn in axonometric projection at a common scale maintained across Figures 4.5 and 4.6; columns follow Table 4.1.

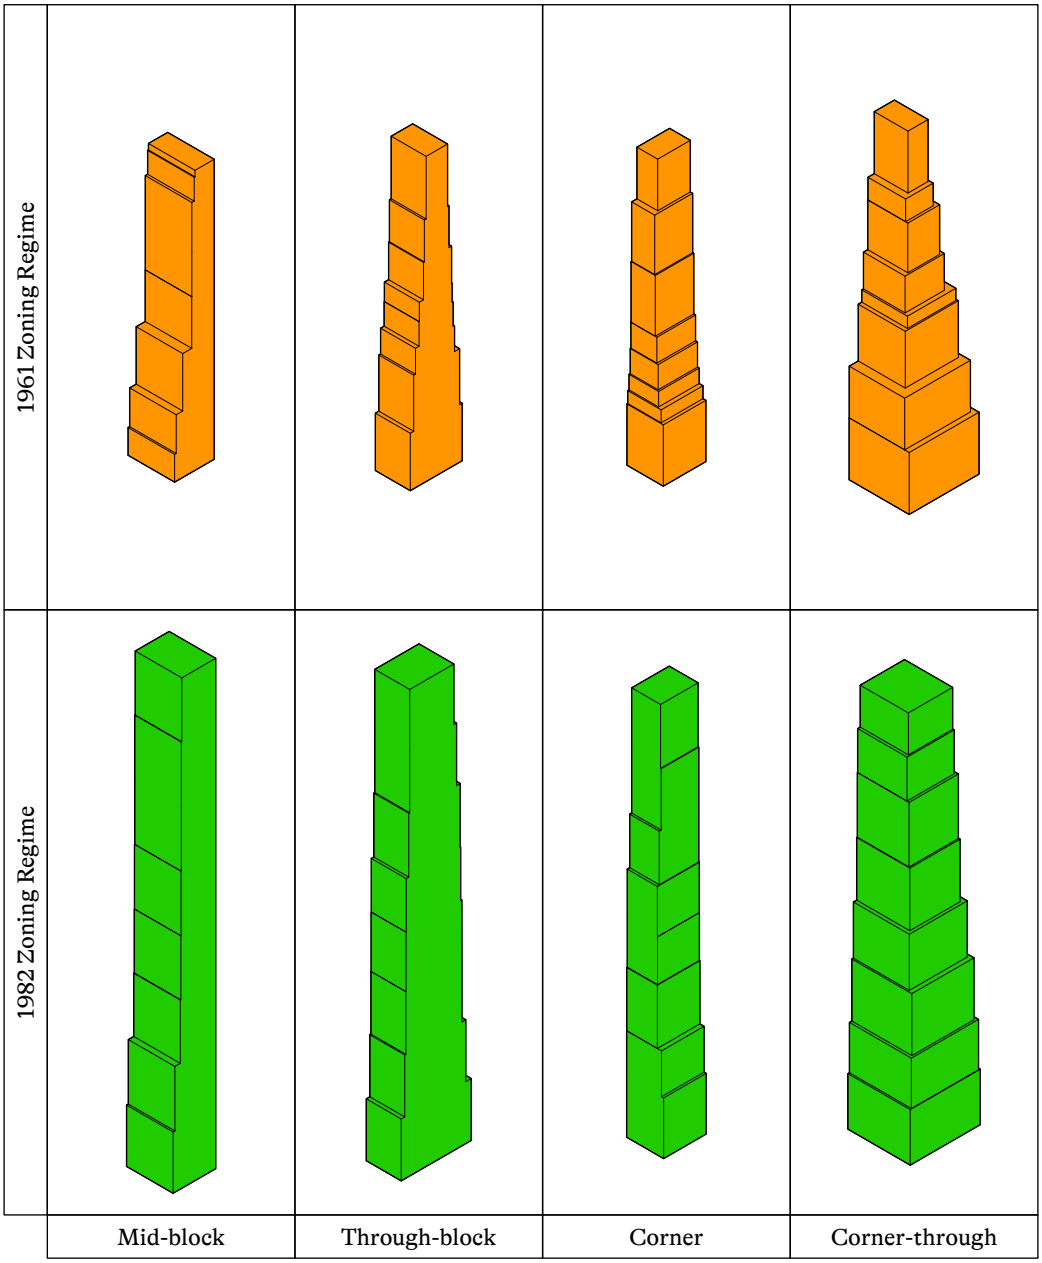


Figure 4.6 — One representative individual per cell, II: the 1961 and 1982 regimes, as Figure 4.5. Read down the pair, the rows trace Table 4.1's ranks in built form: the contraction to 1961 and the recovery under the chart legible in the massing itself.

4.3 H3 — experience: the DEC–SVF relationship

The SVF was measured along the street centreline for the two disjoint samples of §3.6, the off-front interior and the on-front edge, and both were correlated with the DEC score by the Spearman rank coefficient (§3.6). Table 4.2 reports the off-front correlation, the load-bearing test; Table 4.3 the on-front correlation as a secondary check. The off-front coefficient is strong in every one of the sixteen cells, between 0.87 and 0.98; the matrix floor of 0.868 belongs to the 1982 corner.

This result runs counter to H3 as posed. The hypothesis predicted a weak or inconsistent correlation, weaker still within the form-region the 1982 regime opened; the data return the opposite: a strong, consistent correlation across the entire matrix, the 1982 region included. H3 is therefore falsified on this evaluation by the criterion it names. The anti-circularity protocol of §3.6 sharpens the conclusion: the three regimes in which daylight was never the selection objective couple as tightly as the one bred for it, so the coupling is no artefact of selecting for a correlated quantity.

	Mid-block	Through-block	Corner	Corner-through
Pre-1916	0.959	0.941	0.977	0.956
1916	0.977	0.974	0.971	0.982
1961	0.962	0.970	0.942	0.964
1982	0.929	0.925	0.868	0.969

Table 4.2 — Off-front DEC–SVF Spearman correlation, the load-bearing H3 test, across the dominated interior of each evolved population.

	Mid-block	Through-block	Corner	Corner-through
Pre-1916	0.968	0.986	0.906	0.915
1916	0.992	0.989	0.928	0.982
1961	0.950	0.660	0.969	0.995
1982	0.826	0.994	0.897	0.987

Table 4.3 — On-front Spearman correlation, the secondary check on the Pareto edge. It echoes Table 4.2 with two soft cells; the minimum, 0.660 on the 1961 through-block, is read on a front spanning just 14.5 daylight points, where rank correlation is unstable by construction.

The 1982 forms hold no less sky at matched score than the forms of any regime: a common curve through each typology's pooled forms departs from no regime by more than 0.012 of a unit of SVF; nor does the 1982 regime carve out a region of lower-sky compliant form: among legal solutions its minimum centreline SVF is the highest of the four on every lot, 0.011 to 0.022 above the 1916 legal minimum, the unregulated baseline holding the lowest floor throughout. The margin is small but instrument-stable (re-measuring the most exposed cell at the fine ray build moves it by 0.0005, §3.9); it is not controlled against seed variation (§5.3), so the claim rests on the ordering of the gap, not its size. The score threshold prunes precisely the low-sky forms a geometric envelope cannot see.

No systematic departure from tight coupling survives the full matrix. Pooling each typology's four regimes (Figure 4.7), the coupling stands at 0.955 on the mid-block, 0.963 on the through-block and the corner-through alike, and 0.927 on the corner, the one modest softening, traceable to the 1982 corner cell at the matrix floor of 0.868, where two adjacent frontages read a single mass through one coupled sky. An earlier reading of the part-completed matrix had suggested the opposite geography, that the through-block, presenting two independent skies against the chart's one-vantage-per-frontage gaze, would loosen the pair; the completed matrix retires that conjecture. The decoupling the thesis sought appears as a property of no regime, and only faintly, on the corner, as a property of the lot.

Two qualifications bound this verdict. It is drawn from the evolutionary evaluation alone, and H3 is falsified only by a strong correlation in both evaluations; the retrospective cohort supplies the independent replicate. And the two evaluations need not agree: selection rewarded daylight, the market rewards bulk (§5.2), and whether the built stock occupies the minimum-compliant, maximum-bulk corner where score and sky part is the question Figure 4.8 answers.

Across the thirty-six frontages of the retrospective cohort the rank correlation between DEC score and mean centreline SVF is +0.56 ($p < 0.001$); taken building by building the coefficient is +0.53, at fifteen observations near the conventional threshold ($p = 0.041$), a boundary too fragile at this sample size to bear weight in either direction. Against the +0.87 to +0.98 of the evolved populations this is a marked attenuation, and it is decisive: H3 is falsified only by a strong and consistent correlation in both evaluations, and the built record does not supply one. The hypothesis therefore stands.

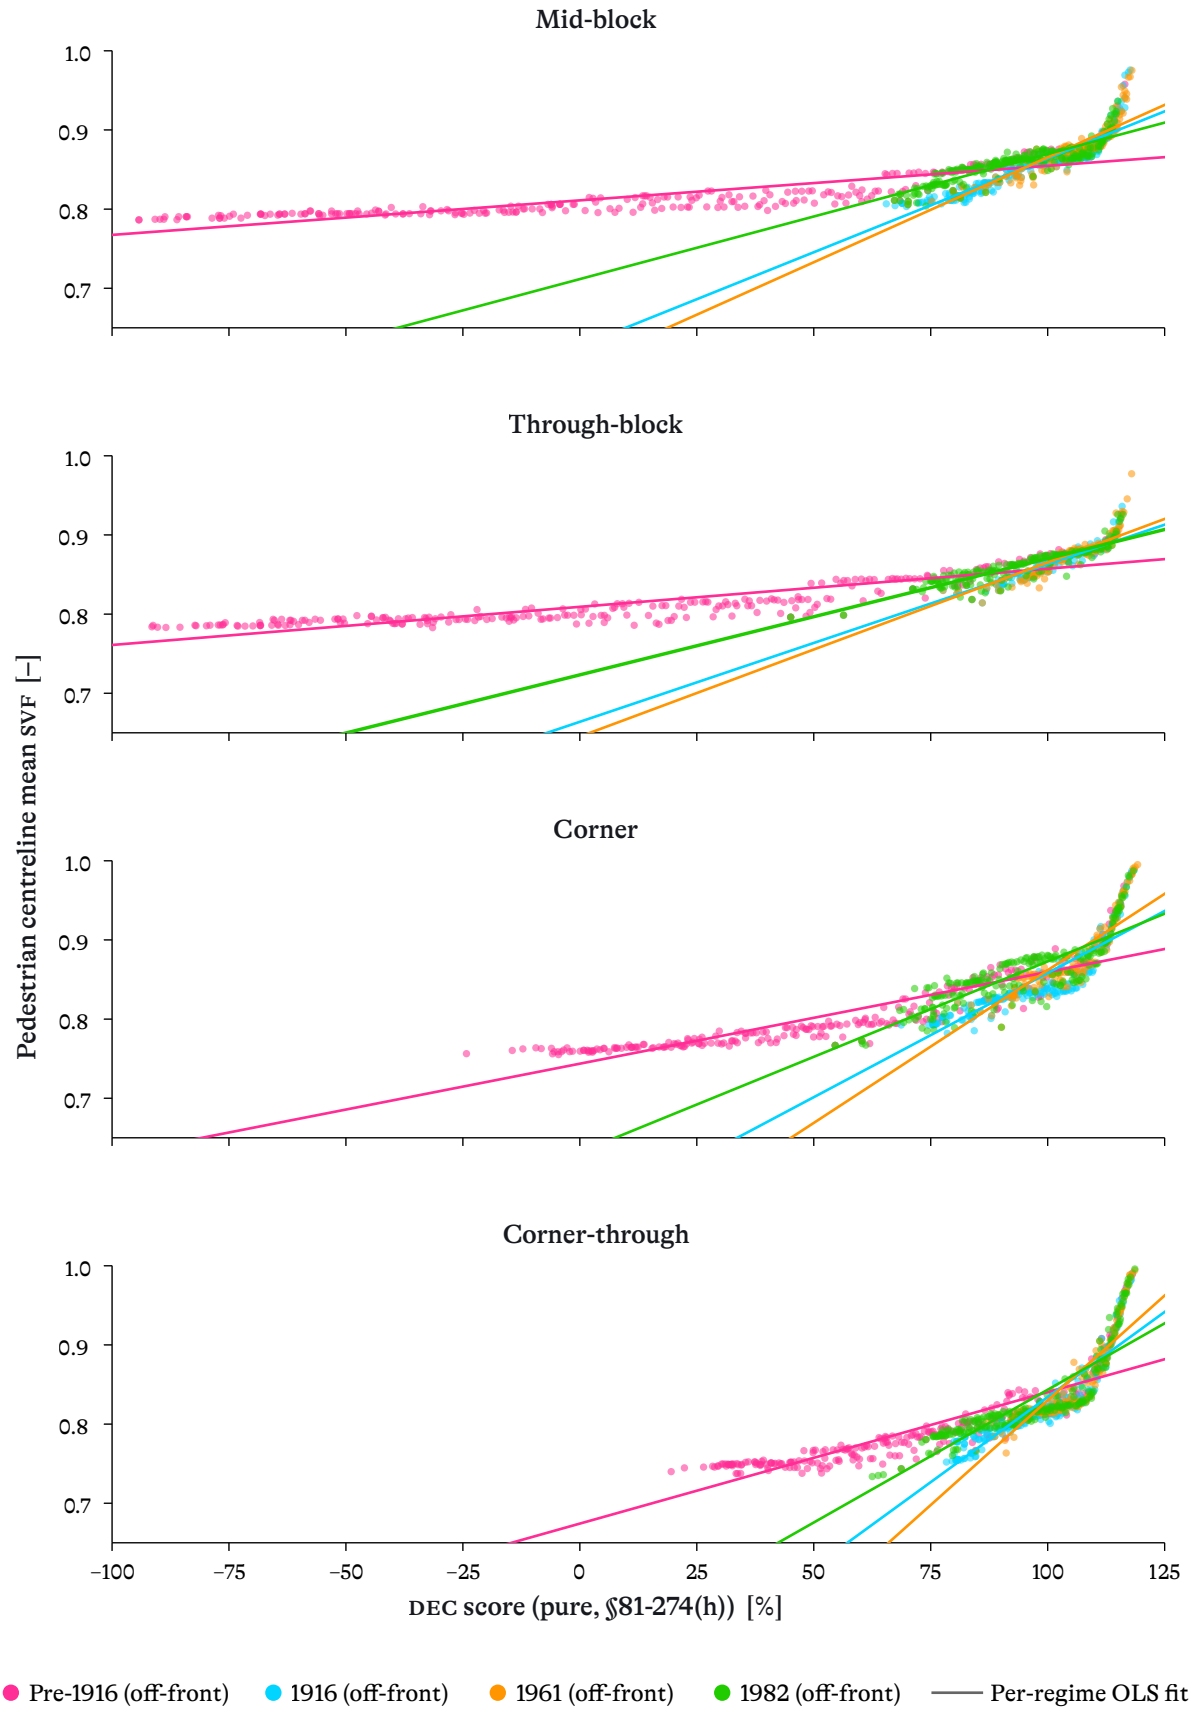


Figure 4.7 — DEC score against centreline SVF, pooled by typology across the four regimes (off-front samples), with per-regime fits; the regimes trace a common curve.

The direction of the divergence is as instructive as its magnitude. The buildings that leave the line leave it conservatively, because the chart penalises a broken setback profile to which the SVF is indifferent: 292 Madison Avenue holds more centreline sky than four of the five 1982 buildings, 0.783, while scoring 69.2%. The metric is stricter than the experience it stands for, not laxer: the pedestrian is not betrayed by the score, but neither is the score a dependable proxy for the sky among the buildings the market actually built.

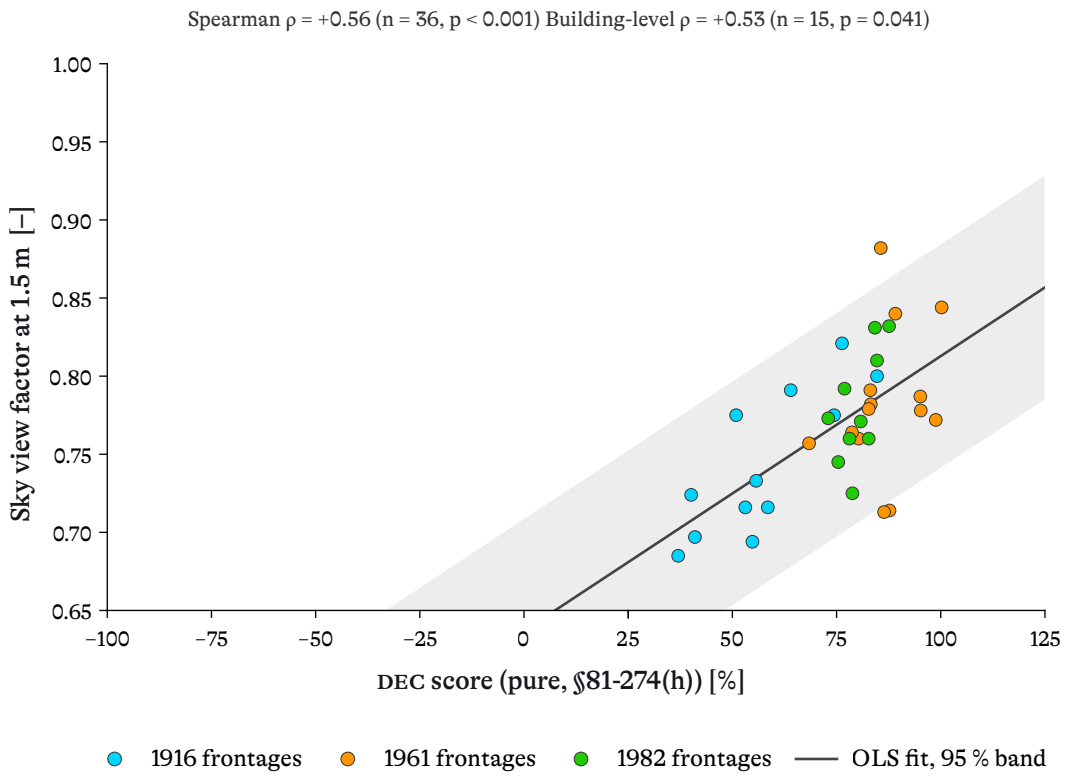


Figure 4.8 — The retrospective cohort: DEC score against mean centreline SVF, set against the relationship of Figure 4.7; one point per frontage and one ring per building. Where the evolved design space couples tightly, the built record couples only moderately.

4.4 Regime findings: the building stock by regime

The systematic screen of §3.8 returned a stock distributed very unevenly across the three post-1916 regimes (Figure 4.9), and the distribution is itself a finding. The 1916 regime, with no floor-area ceiling, left an abundant stock of regular-lot setback towers; the 1982 regime a moderate one; the 1961 regime almost none.

The scarcity is structural rather than accidental. A MAPPLUTTO survey of the 1962–75 office stock shows a sharply bimodal built ratio: a small as-of-right cohort near the district base of fifteen, against a large bonused-and-assembled cohort between eighteen and twenty-seven. The tower-on-plaza the 1961 regime rewarded was worth building only at scale, and scale required assemblage, which carried frontages past the 250 ft limit of the evaluation setup. The canonical towers therefore fall outside the geometry the DEC can describe, and the as-of-right tier the calibration is tested against is the contained-lot minority beneath them. That the instrument calibrated against the 1916 and 1961 stock cannot, by its own construction, score the largest buildings of those regimes is a finding in its own right, and it bounds the retrospective audit accordingly.

The cohort scores make the asymmetry concrete. The 1916 buildings fail the chart not uniformly but locally: their avenue frontages score respectably (84.7% on 42nd Street at the Chanin Building) while their cross-street frontages collapse, to 41.0% on 41st Street and 37.0% at 2 Park Avenue, because a sheer wall on a 60 ft street fills the vantage the chart constructs. The 1961 buildings pass everywhere, most comfortably where the base is most withdrawn: 880 Third Avenue, the smallest building in the cohort, returns 97.7%. What the 1982 regime secured was neither a stricter street than 1916 nor a more generous one than 1961, but a predictable one: the five as-of-right buildings designed against the chart sit within four points of one another, a consistency neither prescriptive regime achieved. The per-vantage origin of these scores is drawn at Figures 4.10–4.12: each cohort's exemplar at its binding vantage.

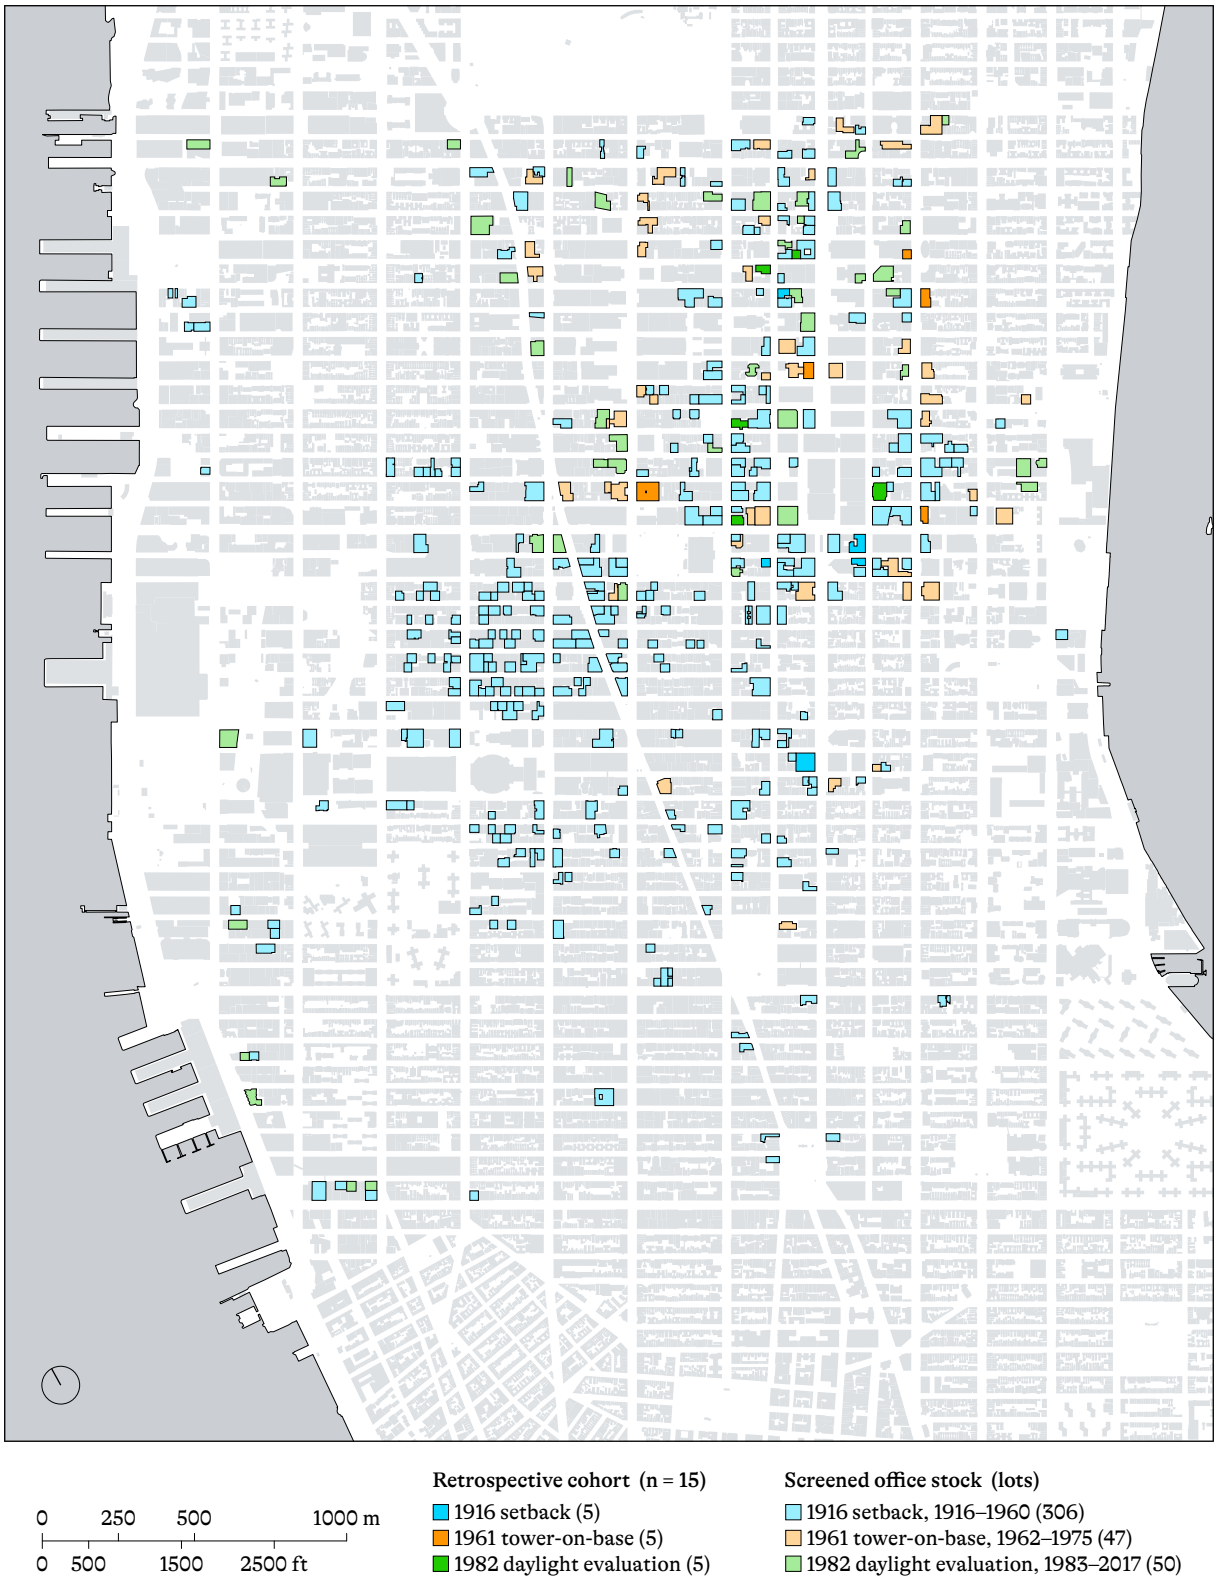


Figure 4.9 — The retrospective cohort in its screened population: MAPPLUTO office lots across community districts 104–106 with a lot area of at least 7,000 sq ft and neither dimension over 250 ft (B.1), classified by completion year into the three regime windows and drawn as building footprints. The screen returns 306 lots under the 1916 regime, 47 under 1961 and 50 under 1982: the asymmetry this section reads. The fifteen cohort buildings of Table B.1 are shown at full saturation; they were drawn from the as-of-right tier within these totals rather than from the totals themselves (Table B.2). Counts are tax lots, and a lot may carry more than one footprint. Avenues vertical: the Manhattan grid runs 29° east of true north.

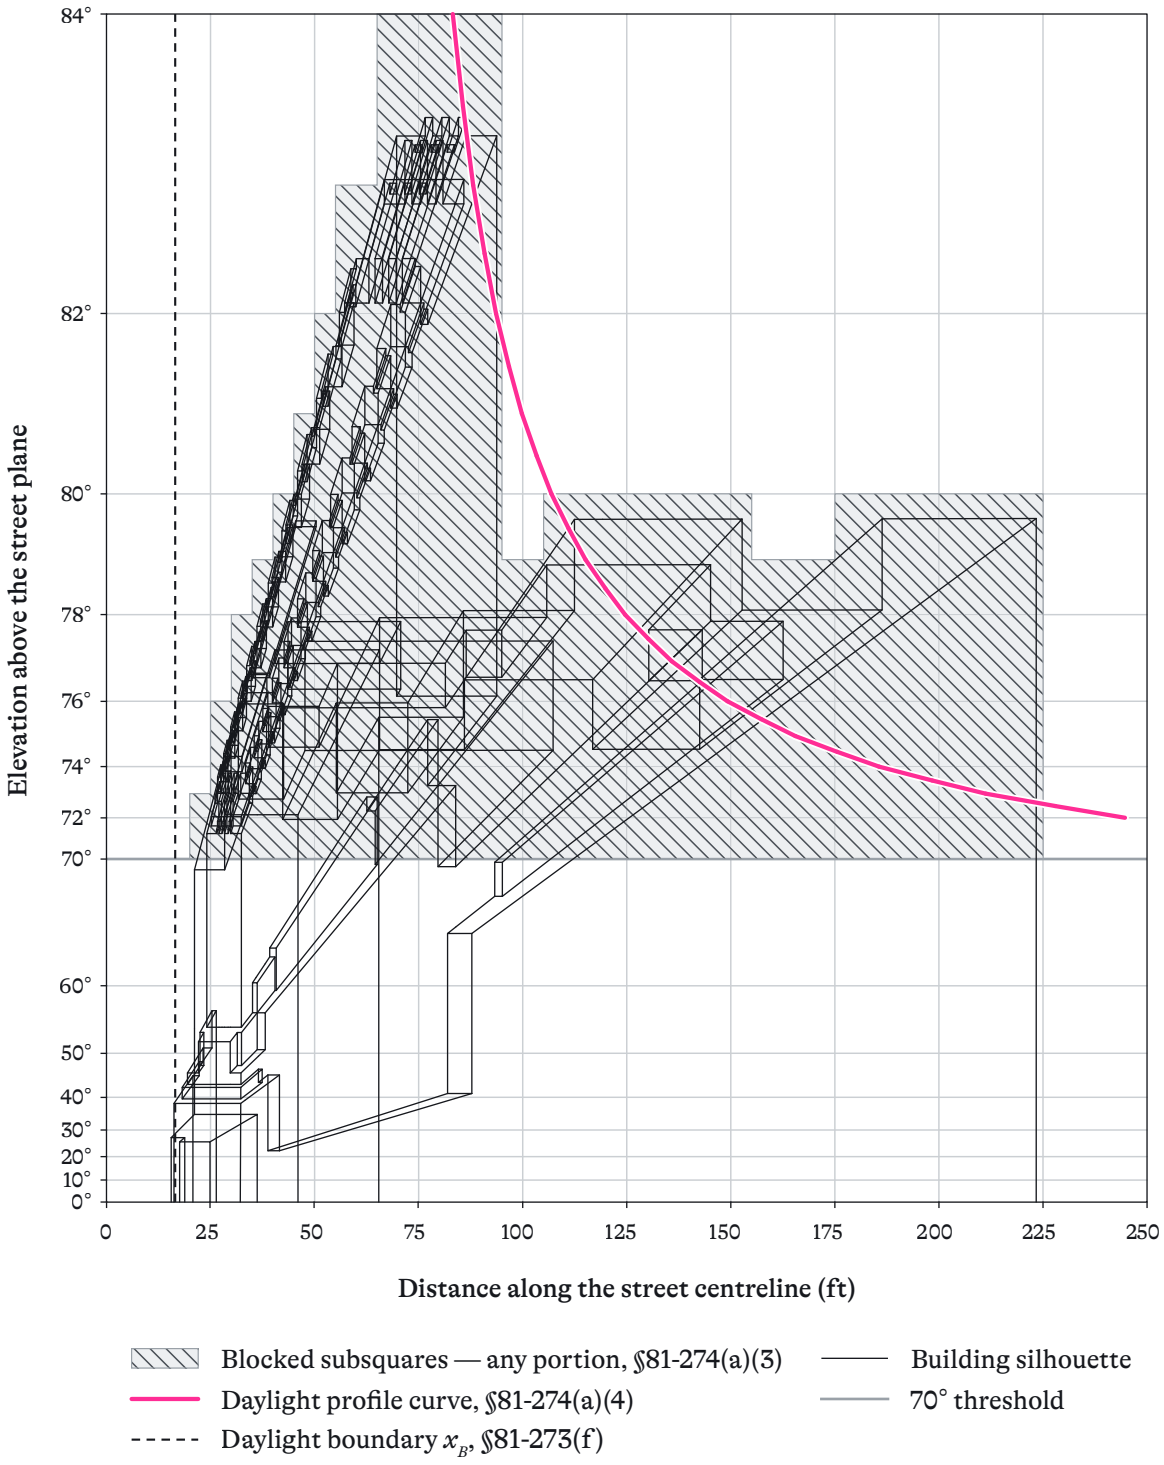


Figure 4.10 — The Chanin Building (1916 cohort) at its binding vantage: the lowest-scoring vantage of its minimum-scoring street, the view the §81-274(h) per-street test reads. On E 41st Street (60 ft), of 94.0 available squares the massing blocks 423 subsquares (−42.3); a below-70° credit of +2.1 and a profile penalty of −16.25, the deepest in the cohort, the setback tower riding far beyond the curve, leave a vantage score of 39.9% and a street score of 41.0%, against the building's overall 58.9%. The shaded subsquares correspond cell-for-cell to the scorer's blocked count.

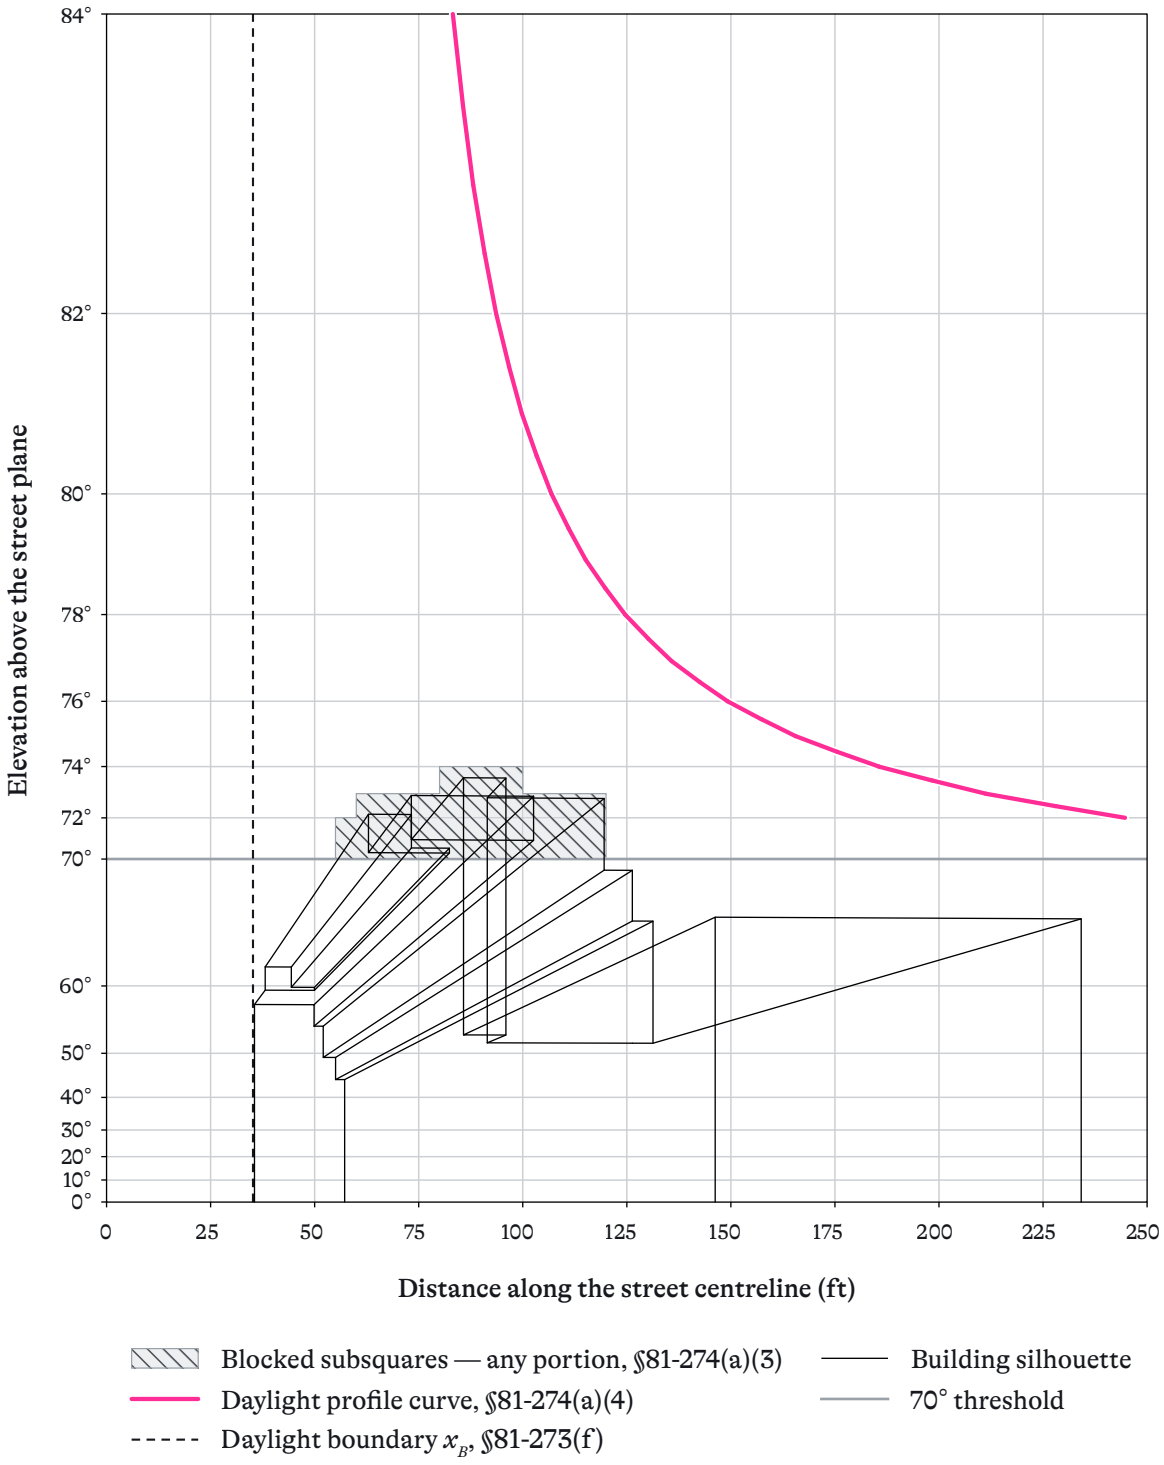


Figure 4.11 — 880 Third Avenue (1961 cohort) at its binding vantage, as Figure 4.10: on E 53rd Street (60 ft), of 86.0 available squares the tower-on-base blocks 42 subsquares (−4.2), with no below-70° credit and no profile penalty, leaving a vantage score of 95.1% against the building's overall 97.7%. A tenth of Figure 4.10's blockage: the contained as-of-right tower holds its mass off the narrow street, and both vantages tie by symmetry.

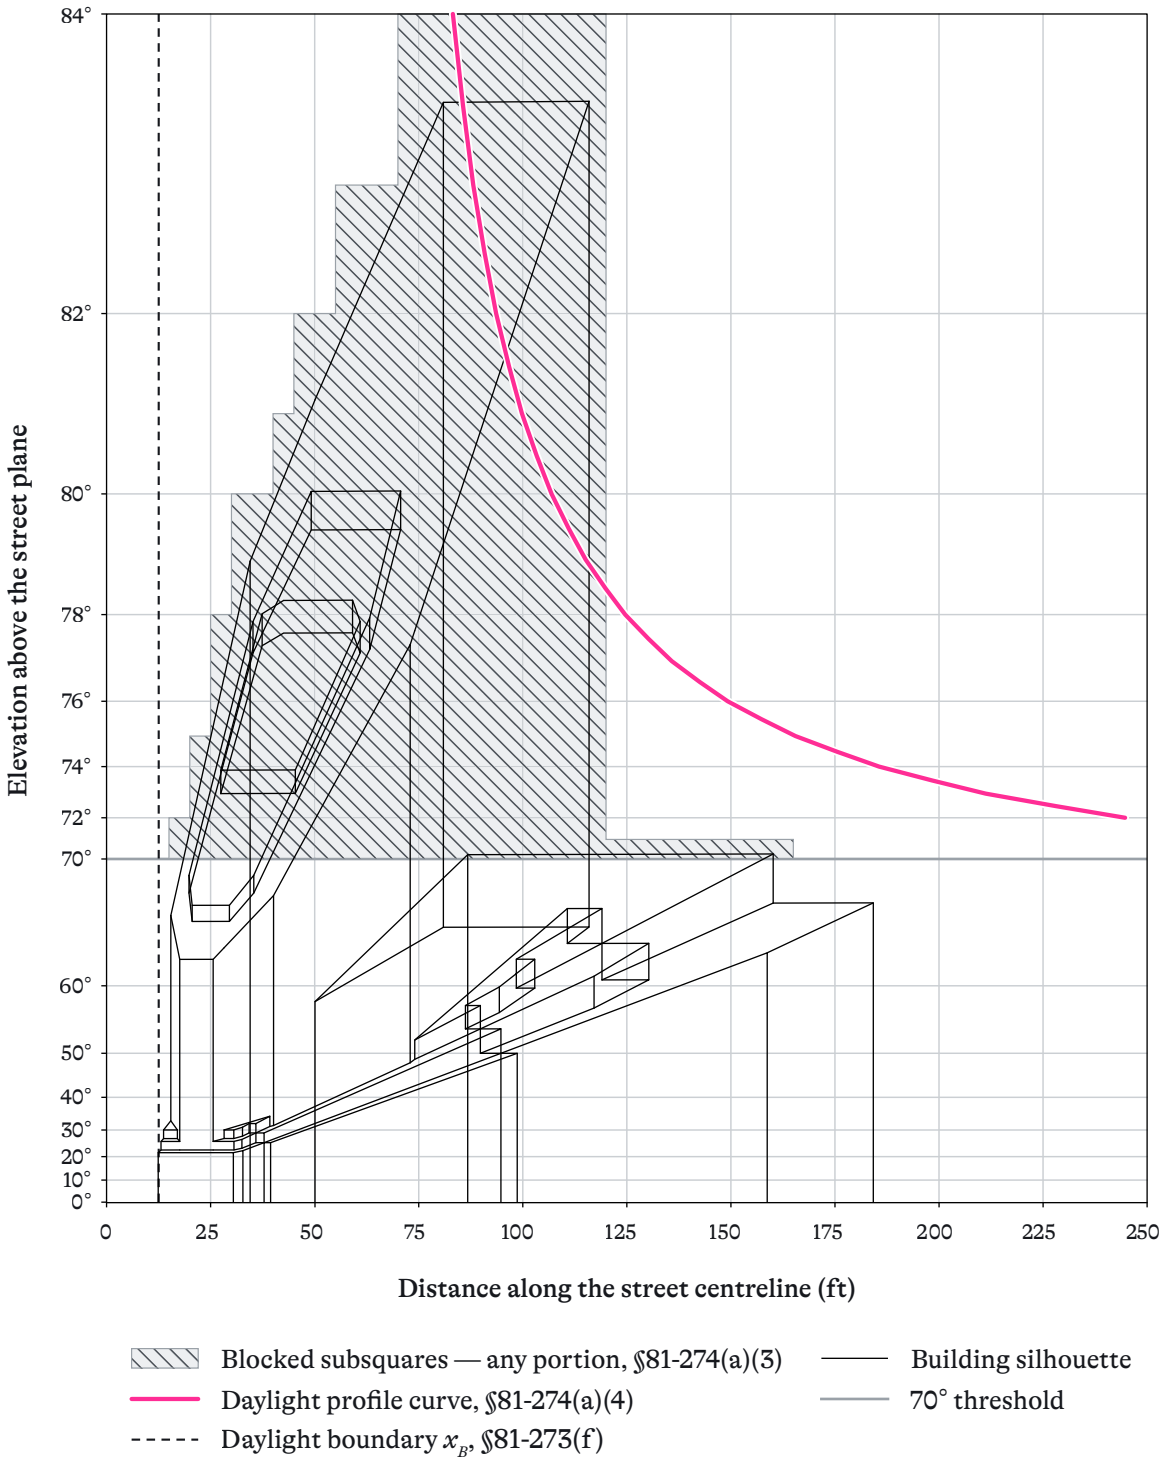


Figure 4.12 — 425 Lexington Avenue (1982 cohort) at its binding vantage, as Figure 4.10: on E 43rd Street (60 ft), of 94.0 available squares the massing blocks 258 subsquares (−25.8); a below-70° credit of +4.2 and a profile penalty of −5.70 (twenty-five subsquares beyond the curve) leave a vantage score of 71.0% and a street score of 73.0%, against the building's overall 78.6%. The shaded subsquares correspond cell-for-cell to the scorer's blocked count.

5. Discussion and Conclusion

5.1 Interpretation

The 1982 change that replaced the envelope with a score had two effects at once, and together they answer the thesis's opening question. It widened what may be built: the chart preserves 0.74 to 0.86 of the unregulated design space, against the envelopes' 0.35 to 0.59, on every lot geometry tested. And the widening cost nothing in the regulation's own currency: at matched DEC score the four regimes deliver, to within 0.012 of a unit of SVF, the same sky along the pedestrian centreline, and the off-front rank coupling never drops below +0.868 in any of the sixteen cells.

Considered together, the two results acquit the chart of the charge implicit in H3, though only within the design space: nowhere in the admitted space can the score be reached without the sky. The chart's threshold proves the sharper guardian of the pedestrian floor: among legal forms the 1982 minimum centreline SVF is the highest of the four regimes on every lot, because a score threshold removes exactly the low-sky forms a geometric envelope cannot see. The audit also identifies two unintended parts of the solution: the decoupled corner of §4.3, where score and sky part company, revealed only by a search of the space of form; and the 250 ft frontier of §81-275, beyond which the chart cannot describe a frontage, met by the audit as a boundary (§4.4).

The design space cannot settle the built record, and the cohort scores (§4.1) do not extend the verdict to the city. Across the fifteen buildings score and sky correlate at +0.56 per frontage and +0.53 per building: a moderate association, not the tight coupling of the evolved populations, short of the strong and consistent correlation H3's falsification requires. What remains is the weaker of H3's two limbs: the claim concerning the 1982 region is refuted on both evaluations, and the departures the built record does show are conservative, though nothing in the rules compelled that margin.

5.2 The confound revisited

The problem posed in §3.6, that an algorithm selecting on the DEC score might manufacture the coupling under test, can be revisited against the complete matrix, and four features of the evidence tell against it. The off-front samples are drawn from the dominated interior and spread across the entire daylight range, not from the optimised edge; the coupling remains at least as tight under the three regimes whose admissible regions owe nothing to the chart's geometry, the prescriptive envelopes coupling at

+0.94 to +0.98 against the chart's own +0.868 to +0.969; and it survives a change in the kind of constraint, holding equally where the decoder clips and where the population trades across the legality line. The fourth feature is the selector's blindness: the SVF was never an objective, calculated only after selection finished, and the search carried no theory of the rule and no intention of any kind. A search that never saw the sky cannot have steered the space towards it, and if the coupling has an author, it is the regulation's own geometry, subject to the channel conceded below.

One channel remains: the objective itself. The selection rewarded daylight; the market rewards bulk. If the minimum-compliant, maximum-bulk form decouples score from sky, the evolved populations will under-represent it, however the samples are stratified. The built record, the one population produced under the market's objective, couples at +0.56 against the evolved populations' +0.87 to +0.98. That gap is consistent with the selection effect no evolutionary control could close: the coupling in the design space was in part a property of what the algorithm sought, not only of what the regimes permitted. A third population was sampled by neither evaluation: forms designed by architects under the regime itself, the competition entries and the unbuilt schemes. The same instrument could score them, asking how forms produced by judgement rather than by optimisation fare under the regulation's own machinery: the most direct test left open.

5.3 Limitations

The ray sampling of the sky-view instrument bounds what its numbers can claim: building-level means converge to within 5×10^{-4} under the retrospective build and to roughly 2×10^{-3} under the evolutionary build, errors demonstrated to move no conclusion, and per-sample values, the weakest quantity, are never used quantitatively in the thesis, the profile figure of §3.5 being drawn from the fine build (C.2).

The verdicts above are bounded by eleven limitations, each documented where it arises. First, the study is computational rather than perceptual: the SVF is a geometric proxy, and no human-subjects evidence shows equal sky reads as equal openness. Second, each cell is a single run from a fixed seed; convergence diagnostics argue the budget sufficed, but replication across seeds was not performed, and Appendix D's fill-forward artefact warns against over-reading any trajectory. Third, the instrument carries documented counting conventions: the canonical any-portion rule of §81-274(a)(3) reads the published worked example at 72.89 against its printed 76.75, a conservative residual C.9 decomposes; cohort scores are read against the anchored expectation of §3.8. Fourth, the volume objective is an open-mesh identity, comparable within a lot but never a measure

of floor area. Fifth, the cohort rests on completion rather than approval dates and on massing verified from imagery, the 1916 sky-exposure multiplier held at its documented constant. Sixth, the DEC cannot describe frontages beyond 250 ft, so the assembled towers of the 1961 regime lie outside the audit and its conclusions (§4.4). Seventh, on a 75 ft street the ZR fixes the vantage at 40 ft, not at the geometric half-width of 37.5; the cohort is scored at the regulatory vantage, the matrix is not, and the deviation is identical across regimes within each cell, so §4.2's ratios and §4.3's correlations are unaffected (B.5). Eighth, the as-of-right gate removes the buildings most likely to occupy the decoupled corner (§4.3): the maximum-bulk form is built on merged, bonused lots, and the audit excludes it by construction (B.9). Ninth, both instruments read the subject building alone, with no surrounding massing: the chart's own convention (§3.5), and the condition under which the two measures are commensurable. The centreline SVF values reported here, 0.70 to 0.82, are therefore isolated-building quantities, not the sky a pedestrian sees on a built-up street; the comparative results are unaffected because the convention is uniform, but the absolute figures are not to be read as measured street conditions. Tenth, the test case is a favourable one: sixteen cells on the Manhattan grid, regular, typologically well defined and geometrically rigid, so the sampling meets none of the irregularity of older street networks. The numeric thresholds are not claimed to transfer elsewhere; what transfers is the protocol: permissiveness measured against an unregulated baseline, coupling between score and sky, and calibration re-derived from local stock, each defined on the instruments, not the geometry. A second city would be the test of that transfer. Eleventh, the evolved space is the space of the expressible, not of the intended: an eight-segment orthogonal grammar searched under a daylight-and-volume objective, with no account of what an architect would want to build. Forms outside the grammar, or unrewarded by the objective, are absent by construction, the extreme corners of the admitted space products of the rule, not of design intention.

Three choices would differ with hindsight: context massing admitted in a second SVF pass, at the cost of the chart-parity the thesis exploits (§3.5); seeds replicated into a distribution over runs; and the merged-lot maximum-bulk form modelled into the audit rather than excluded from it (B.9).

None of this argues against the method. The space of what a regulation permits cannot be enumerated by hand; evolutionary search maps it precisely because it is theory-blind, the appropriate instrument for auditing a rule whose intent is itself under examination. Convergence diagnostics, the stability of every headline ordering, and the reproducibility of the compiled instrument argue that the search, not the noise, produced the findings; the single-seed runs remain the honest boundary of that claim.

5.4 Contribution and future work

The contribution is a way of reading regulation, an instrument, and the findings the two together produced. The method treats a zoning regime not as a restriction on a building but as a generator of design space, measured as the HV its Pareto front dominates against a common unregulated baseline. The instrument is an open, validated implementation of the DEC, anchored to the regulation's published worked example, declaring the counting convention the chart leaves unstated, and paired with a pedestrian sky-view measure the chart itself cannot take. Together they yielded the permissiveness order pre-1916 > 1982 > 1916 > 1961, consistent across the four lot typologies; the discovery that the prescriptive regimes respond to lot geometry in nearly opposite directions, while the chart redistributed rather than flattened the regulation's geographic incidence; the tight, regime-indifferent coupling of score to street-level sky; and, from the built record's own distribution, a fourth: the 1961 regime, whose canonical towers the ruling instrument cannot even describe.

Beyond Midtown, the technique carries its strongest claim: no comparable end-to-end audit of a zoning instrument in force appears in the literature reviewed here, and a performance-based rule should carry such a harness as a condition of adoption, regulation as code. Any rule that admits form by score inherits the two conditions demonstrated here: its arithmetic must be specified to the counting convention, or its scores are not comparable; and its calibration is an empirical claim, open to re-audit as the stock it was derived from changes. The chart's authors calibrated from the buildings already standing, before the space a rule would admit could be searched; a proposed revision could now enter the harness before enactment rather than be audited decades after. The instrument is built for that scrutiny: a pure function with declared conventions, released in the digital annex, so every score here can be recomputed and challenged under the rule that produced it. The object is current law: the December 2024 citywide amendments left the daylight machinery untouched (§2.5), so what H1 tests is the calibration of the rule now in force, not a closed historical episode.

The harness does not replace the planner's judgement. The score was adopted in 1982 partly because discretionary departures had eroded the envelope (§2.3), and narrowed the scope for discretion without eliminating it: special permits and floor-area bonuses remain in force, and the study's own gate set the bonused stock aside, an as-of-right cohort of five at the district base of fifteen against a bonused-and-assembled majority carrying eighteen to twenty-seven (§4.4; Table B.2). The instrument provides an as-of-right baseline against which departures can be judged: a special permit can now be

asked what it costs in street-level sky, in the rule's own units. The counting-convention finding is a reminder that the chart leaves room for interpretation; the computational case informs the discretion that remains, it does not pretend to remove it.

The nearest task is the compliance route of each 1982 building, on zoning diagrams not yet digitised. Beyond it, three extensions follow: a human-subjects study of whether equal SVF is perceived as equal openness, converting the geometric proxy into a perceptual claim; the post-1982 amendments and the 2017 East Midtown rezoning, further regimes for the same sixteen-cell harness; and a third from the residual C.9 decomposes. Because the chart charges blockage by subsquare, its score is an area integral truncated to a cell count; the same rasteriser can return the exact area-weighted blockage, bounding the legislated grid's resolution error.

The furthest extension is infrastructural: the regime modules encode a zoning regime as executable code (Appendices A and E), and the same encoding could be written for other cities, grids less regular, typologies less settled, street widths less canonical, where the couplings and ratios demonstrated here would be the first quantities re-tested and the harness could accompany a regulation from drafting. Running the validation and the sixteen-cell orderings through a second engine and optimiser would test both instrument and findings. The thesis set out to catch a regulation preserving its metric while losing its street. In the space of what may be built, the chart keeps its promise; in the city it produced, the promise holds only loosely.

References

Works cited

Caldas, L.G. and Norford, L.K. (2002). 'A design optimization tool based on a genetic algorithm', *Automation in Construction*, 11(2), pp. 173–184.

Deb, K. (2000). 'An efficient constraint handling method for genetic algorithms', *Computer Methods in Applied Mechanics and Engineering*, 186(2–4), pp. 311–338.

Deb, K., Pratap, A., Agarwal, S. and Meyarivan, T. (2002). 'A fast and elitist multiobjective genetic algorithm: NSGA-II', *IEEE Transactions on Evolutionary Computation*, 6(2), pp. 182–197.

Ferriss, H. (1929). *The Metropolis of Tomorrow*. New York: Ives Washburn.

Fishhaut, E.J., Weingarten, N.H. and Neiger, C.J. (1984). 'An alternative to the Daylight Evaluation Chart', in Wexler, J. (ed.) *CAD84: 6th International Conference and Exhibition on Computers in Design Engineering*. London: Butterworths, pp. 182–192.

Guerreiro, A.P., Fonseca, C.M. and Paquete, L. (2021). 'The hypervolume indicator: computational problems and algorithms', *ACM Computing Surveys*, 54(6), Article 119, pp. 1–42.

Jacobs, J. (1961). *The Death and Life of Great American Cities*. New York: Random House.

Kwartler, M. and Masters, R. (1984). 'Daylight as a zoning device for Midtown', *Energy and Buildings*, 6(2), pp. 173–189. doi: 10.1016/0378-7788(84)90072-0.

Liu, K., Xu, X., Huang, W., Zhang, R., Kong, L. and Wang, X. (2023). 'A multi-objective optimization framework for designing urban block forms considering daylight, energy consumption, and photovoltaic energy potential', *Building and Environment*, 242, Article 110585. doi:10.1016/j.buildenv.2023.110585.

Oke, T.R. (1981). 'Canyon geometry and the nocturnal urban heat island: comparison of scale model and field observations', *Journal of Climatology*, 1(3), pp. 237–254.

Ratti, C., Raydan, D. and Steemers, K. (2003). 'Building form and environmental performance: archetypes, analysis and an arid climate', *Energy and Buildings*, 35(1), pp. 49–59.

Wang, L., Janssen, P. and Ji, G. (2024). 'Optimization-based design exploration of building massing typologies — EvoMass and a typology-oriented computational design optimization method for early-stage performance-based building massing design', *Frontiers of Architectural Research*, 13(6), pp. 1400–1422.

Willis, C. (1995). *Form Follows Finance: Skyscrapers and Skylines in New York and Chicago*. New York: Princeton Architectural Press.

Primary regulatory sources

City of New York. Zoning Resolution of the City of New York. Article VIII, Chapter 1: Special Midtown District. §81-25, General Provisions Relating to Height and Setback of Buildings; §81-253, Special Provisions for the East Midtown, Theater, Fifth Avenue, Penn Center and Preservation Subdistricts: establishes that §§81-26 and 81-27 are supplemented and modified within those subdistricts.

City of New York. Zoning Resolution of the City of New York. §81-26, Height and Setback Regulations — Daylight Compensation: the prescriptive alternative to daylight evaluation, used in labelling the 1982 cohort's compliance routes.

City of New York. Zoning Resolution of the City of New York. §81-27, Alternate Height and Setback Regulations — Daylight Evaluation: the performance route implemented by this thesis.

City of New York. Zoning Resolution of the City of New York. §81-272, Features of the Daylight Evaluation Chart. Last amended 2 February 2011. Specifies the 250-foot vantage point, the three charts by street width (60-foot; 75- or 80-foot; 100-foot and wider), the daylight squares and subsquares, the seventy-degree line, and the profile curve.

City of New York. Zoning Resolution of the City of New York. §81-273, Rules for Plotting Buildings on the Daylight Evaluation Chart. Last amended 2 February 2011. Contains the worked example used for validation in Appendix C.

City of New York. Zoning Resolution of the City of New York. §81-274, Rules for Determining the Daylight Evaluation Score. Last amended 2 February 2011. Specifies blockage counting, the below-seventy-degree credit, the profile penalty weights, the aggregation of vantage and street scores, and the passing thresholds of 75 per cent overall and 66 per cent per street frontage, with the Resolution's statement of their derivation from the as-of-right 1916 and 1961 stock.

City of New York. Zoning Resolution of the City of New York. §81-275, Special Conditions. Governs lots outside the standard evaluation setup, including frontages exceeding 250 feet; used as a cohort-selection gate in Section 3.8 and otherwise outside scope.

City of New York. Zoning Resolution of the City of New York. §81-66 (East Midtown Subdistrict, added 2017), including §81-663 and Chart 4 (Daylight Evaluation Diagram – Park Avenue), used in lieu of the 100-foot chart for qualifying East Midtown sites with Park Avenue frontage. Acknowledged in Chapter 2; excluded from scope.

City of New York. Zoning Resolution of the City of New York. Appendix A, Map 1: Special Midtown District and Subdistricts. Establishes the boundaries of the Fifth Avenue and Theater Subdistricts referenced in Appendix B.8.

New York City Department of City Planning. Special Midtown District daylight evaluation charts [official chart graphics]. Available at: https://www.nyc.gov/assets/planning/download/pdf/zoning/graphics-charts/special-midtown/. Source graphics for the profile-curve digitisation described in Chapter 3.

Historical regulatory instruments

City of New York (1916). Building Zone Resolution. Adopted by the Board of Estimate and Apportionment, 25 July 1916. The setback regime encoded as the 1916 constraint module; consulted in the secondary accounts of Willis (1995) and Kwartler and Masters (1984).

City of New York (1961). Zoning Resolution. Adopted by the City Planning Commission and Board of Estimate, effective 15 December 1961. The FAR encoded as the 1961 constraint module; tower, coverage, and setback provisions consulted in the current Resolution's Article III, Chapter 3 lineage and the account of Kwartler and Masters (1984).

Software and data

Makki, M., Showkatbakhsh, M. and Song, Y. (2024). *Wallacei*: an evolutionary and analytic engine for Grasshopper 3D. Available at: https://www.wallacei.com (Accessed: 18 August 2026).

New York City Department of City Planning (2026). *MAPPLUTO*. New York: NYC Department of City Planning. Available at: https://www.nyc.gov/site/planning/data-maps/open-data.page (Accessed: 18 August 2026).

Robert McNeel & Associates (2024). *Rhinoceros 3D* and *Grasshopper*. Seattle: Robert McNeel & Associates.

Scoring the Sky: Digital Annex (2026). DECScorer, DECProjection, CenterlineSVF, the four regime modules, PhenotypeBuilder, HypervolumeCalc and test rigs, with run records for the sixteen cells and the retrospective cohort tables [C# components, Python scripts and data]. Submitted with this dissertation (BARC0141, University College London).

Appendices

Appendix A: Regulatory Regime Reconstructions

Each regime is realised as a Grasshopper C# constraint module (Appendix E) that reads a candidate massing and returns a legality factor in the unit interval: identically one under the prescriptive regimes, whose envelopes the decoder enforces by clipping before any scoring, and under the 1982 regime a soft factor reduced in proportion to the shortfall against its thresholds (§3.7). This appendix records the specific parameter values those modules use, with the section-level citations the main text gives only in principle. Two honesty notes bound what follows. The 1916 and 1961 modules use calibrated working values drawn from Kwartler and Masters (1984) and the surviving resolution texts rather than a live geographic layer, since the 1916 height districts were abolished in 1961 and no longer exist as mappable data; where a value could not be pinned to a primary source it is stated as a calibrated assumption. And under the 1982 regime legality is graded rather than gated: a non-compliant form remains in the population with its objectives scaled by its legality score, a multiplicative penalty of the family reviewed by Deb (2000), used in place of the hard constrained-domination rule that paper proposes and NSGA-II adopts (Deb et al., 2002), precisely so that gradation is preserved; and the reported Pareto front of each regime is the front of legal forms only.

A.1 Pre–1916 baseline

The pre-1916 module imposes no constraint. It applies no height limit, no setback requirement, no coverage limit, and no floor-area ratio, and returns a legality score of exactly one for every massing. It corresponds to the leftmost condition in Kwartler and Masters' comparison of regimes, unlimited height at full lot coverage: the unregulated envelope against which the three later regimes are read, and the reason the pre-1916 column of the design-space results is the upper bound of admissible form. Its module exists for compositional symmetry: every regime has one, and this one happens to constrain nothing.

A.2 1916 setback regime

The 1916 module tests three rules of the setback resolution, per street frontage, and returns the worst-case compliance ratio across them. The street wall may rise vertically at the lot line to twice the street width, the street-wall ratio being set to 2.0, the value Kwartler and Masters give for the prevalent Midtown Height District 2 (pp. 175, 180);

above the street wall the mass must fall beneath a sky exposure plane rising at four vertical to one horizontal, measured from the street centreline rather than the lot line. Above that plane a tower of unlimited height is admitted provided it covers no more than a quarter of the lot area and stands at least seventy-five feet (22.86 m) from each street centreline. On a mid-block lot the plane applies to the single frontage; on corner, through-block, and corner-through lots it applies to each street independently. The street-wall multiplier and the plane slope are calibrated working values taken from the 1916 resolution text and authoritative secondary accounts (Willis, 1995), not a live geographic layer, since the 1916 height districts were abolished by the 1961 resolution. The Fifth Avenue corridor's steeper multiplier and the relaxation on streets wider than one hundred feet are documented exceptions the module does not separately encode, the study's lots falling under the prevalent district. Because the street wall stands at twice the street width and the plane rises four to one, the two constructions coincide: the allowed height at any point is four times its distance from the street centreline.

A.3 1961 tower-on-base regime

The 1961 module tests the floor-area ratio, coverage, and setback rules of the commercial Midtown district. Total floor area, computed by discretising the mass into storeys at a representative floor height of 3.66 m (12 ft), must not exceed fifteen times the lot area, the as-of-right commercial ceiling (Kwartler and Masters, 1984, Figures 4–7; Figure A.1). Above the street wall the tower footprint must not exceed 40% of the lot and total coverage must not exceed 55%, and the mass must set back from the street line, deeper on narrow streets (15 ft) than on wide ones (10 ft), with an alternate front-setback formulation admitted to 35 ft and the corner-lot relaxations of the resolution's tower rules. The module also grants the resolution's plaza floor-area bonus: floor area is credited in exchange for ground-level plaza, raising the effective ratio to a capped maximum of 20% above the base. It is this bonus, not any coverage rule, that rewards a recessed footprint; ground coverage is reported but, consistent with the commercial districts' absence of a ground-coverage cap, is not itself constrained, so the plaza emerges from the incentive rather than from a coded rule. As with the 1916 module, the output is a graded compliance score. The floor-area ceiling and the bonus cap are calibrated to the resolution and to Kwartler and Masters' figures, with the plaza bonus modelled as a proxy for the resolution's dimensioned plaza provisions.

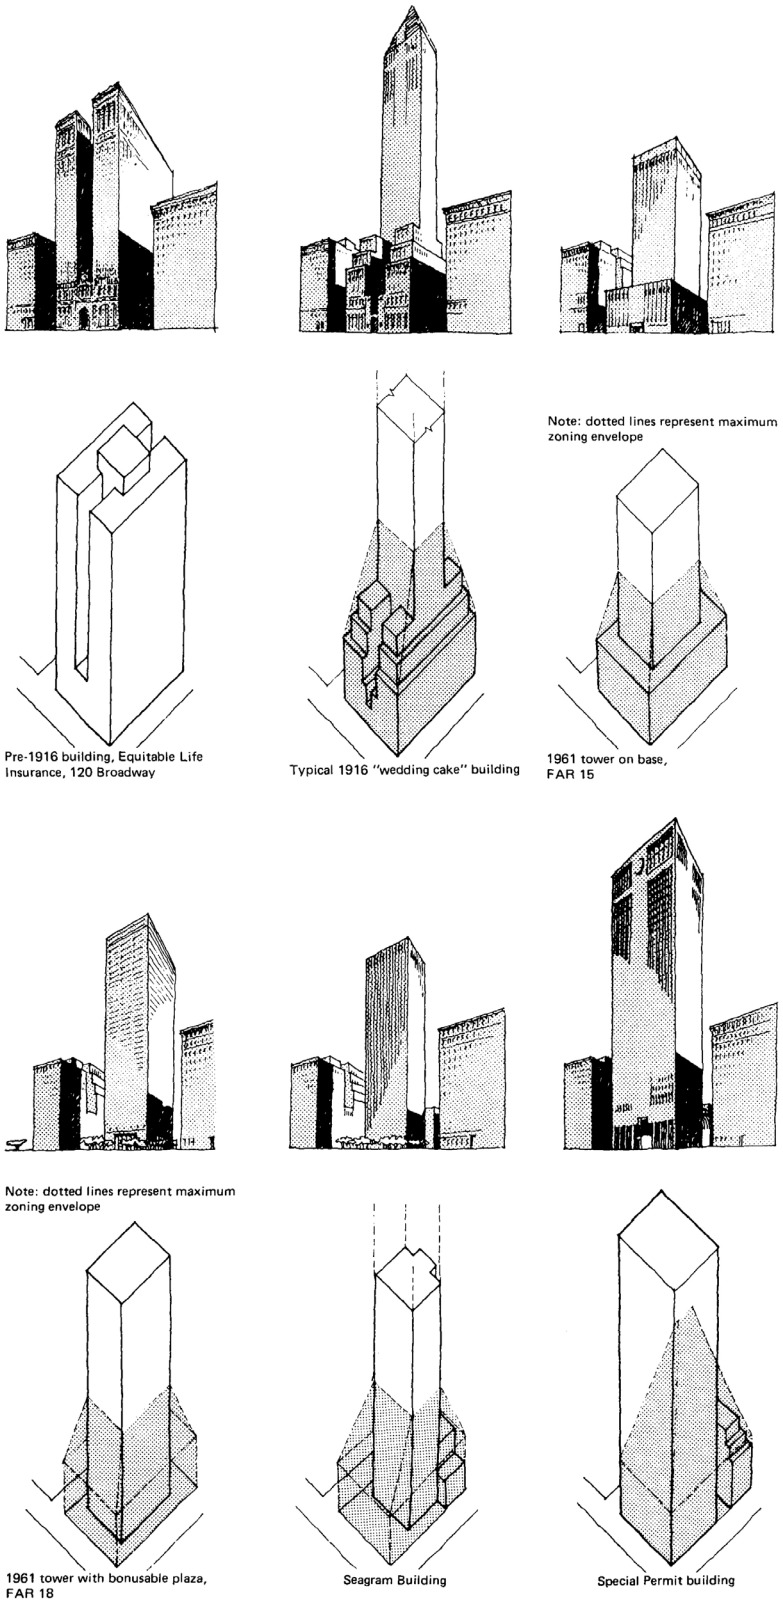


Figure A.1 — The 1961 type at its two floor-area tiers: the as-of-right tower-on-base at FAR 15, and the plaza-bonused variant at FAR 18. Source of the floor-area ceiling and the bonus cap encoded in the 1961 module. Reproduced from Kwartler and Masters (1984, Figures 4-7).

A.4 1982 DEC regime

The 1982 module is a performance test: it wraps the DEC scorer (Appendix C) and admits any massing that meets the two statutory thresholds: an overall score of at least 75% (§81-274(i)) and a score of at least 66% on every street. It imposes no setback envelope, no floor-area ratio, no coverage limit, and no height limit; the Chart alone is the legality test. The scorer builds its vantage points on the street centreline two hundred and fifty feet from the far side lot line and trims the evaluation at a boundary depth of one hundred feet, the standard setup of §81-273(f), which takes the lesser of the centreline distance and one hundred feet. Legality is graded rather than binary: a compliant massing scores one, and a non-compliant massing scores one minus the larger of its two shortfalls, the overall shortfall below seventy-five and the worst per-street shortfall below sixty-six, each as a fraction of its threshold. The route modelled is the as-of-right daylight-compliance path of §81-27; the discretionary alternatives are out of scope, as recorded in A.5.

A.5 Scoped-out provisions

Several provisions that operate alongside the as-of-right rules are deliberately excluded, so that the study compares the base envelope each regime imposes rather than the discretionary departures from it. Excluded are the discretionary and bonus mechanisms, namely special permits and floor-area bonuses obtained through review rather than as of right, whose departures from the as-of-right envelope motivated the 1982 reform; the special-frontage conditions of §81-275, which apply where a single frontage exceeds two hundred and fifty feet, beyond the geometry of the standard evaluation setup; the reflectance modifications of §81-276; and all post-1982 amendments, including the 2017 East Midtown rezoning, which lies beyond the study's cut-off. Each is named here as explicitly outside the scope committed to in §2.5.

A.6 Encoding the regimes as constraint modules

Each of the four regimes is encoded as a constraint module that determines which building forms are admissible. The pre-1916 baseline applies no height, setback, or coverage limit, representing the unregulated condition against which the later regimes are read. The 1916 module applies the street-wall height, sky exposure plane, and tower-coverage rules of the setback resolution. The 1961 module applies a floor-area-ratio limit, coverage limits, and the tower-on-base envelope, including the required setbacks above the street wall. The 1982 module applies the DEC as a performance constraint, admitting

any massing that meets the 75% overall and 66% per-frontage thresholds. Massing is generated on a building footprint that is distinct from the lot boundary, so that the recessed, plaza-fronted ground condition the 1961 regime was written to encourage is expressible, and judged, architecturally: a recessed footprint is legal but volumetrically wasteful under the 1916 envelope, whose sky exposure plane rises from the street line regardless of where the building stands, and rewarded under the 1961 regime, which grants bonus floor area in exchange for ground-level plaza. The street wall of the first regime and the plaza of the second thus both emerge from the optimisation rather than being prescribed to it.

A.7 Objective framing and the placement of the legality penalty

The asymmetry in how legality enters the fitness, enforced by construction in the prescriptive regimes, applied as a soft penalty in 1982, is set out in §3.7; what follows is the objective framing that accompanies it. Both objectives are framed for minimisation, the convention the algorithm requires, by negating the maximised daylight and volume quantities and computing the HV of every regime against one fixed reference point. In the performance regime, the daylight objective alone carries no legality multiplier: daylight being already the quantity selected, multiplying it by a penalty derived from itself would double-weight the same gradient. The penalty is applied instead to the volume objective, so that built bulk is steered toward legal daylight while the daylight score remains comparable, as "daylight achieved", across all four regimes.

Appendix B: The Retrospective Cohort

The retrospective cohort is fifteen Midtown buildings, five per regime, chosen by systematic screen rather than reputation. This appendix records the screen, the buildings returned, the setback check on the 1916 five, and the population counts by regime. The era's best-known towers are absent, failing the frontage gate or approved by special permit: Empire State, Chrysler, Seagram, General Motors and CBS.

B.1 MapPLUTO screening criteria

The screen is applied to the New York City MAPPLUTO tax-lot layer across the three Midtown community districts (104, 105, and 106). A lot enters the candidate pool if its building class is office (the O-family classes), its lot area is at least seven thousand square feet, and neither lot dimension exceeds two hundred and fifty feet, the geometric limit of the standard evaluation setup, beyond which the special conditions of §81-275 apply. The regime is assigned by year of construction: 1916 to 1960 for the setback regime, 1962 to 1975 for the tower-on-base regime, and 1983 to 2017 for the daylight regime, the last bounded above by the 2017 East Midtown rezoning. The screen reads mappluto's tabular frontage, depth and area; the modelled lot lines and the frontage lengths that weight the scores follow its lot polygons, which run a few feet longer.

B.2 The fifteen buildings

The screen of B.1 returns the population; the cohort is drawn from a tier within it. For the 1961 and 1982 regimes that tier is the screened lots of fifteen storeys or more whose built floor-area ratio lies in a band about the district base, 12 to 16.5 under 1961 and 12 to 21 under 1982, the bonused tier being everything above it (Table B.2). The 1916 regime has no floor-area ceiling, so its tier is defined by height and regularity instead: of the 306 screened lots, 214 rise to fifteen storeys or more, and 95 of those carry MAPPLUTO's regular-lot flag (IrrLotCode = N), reported in Table B.1's "Regular" column. Neither tier filters on lot typology, and the 1916 tier includes interior lots. From each tier the five buildings were those passing the fourth gate of §3.8 and, for 1916, the setback check of B.3; the Chanin Building, the pilot, was admitted before the screen and is irregular by that flag. The cohort therefore holds no 1916 interior lot because of what could be modelled and confirmed, not what the screen returned; §4.1 records what that leaves untested. Table B.1 lists the fifteen with their tax-lot identifiers; lot dimensions are frontage by depth in feet (the Chanin Building's three frontages given in full), floor-area ratios the as-built values from the tax-lot record.

Regime	Building	Compl.	ZD	Lot (ft)	FAR	Regular	BBL
1916	Chanin Building	1929	C5-3/2.5	125/175/197.5	29.10	No	1012960014
1916	364 Lexington Ave	1930	C5-3	98.75×150	16.65	Yes	1012950058
1916	2 Park Ave	1928	C5-2	197.5×205	21.06	Yes	1008620029
1916	485 Madison Ave	1929	C5-3	100.4×124.5	19.57	Yes	1012870052
1916	292 Madison Ave	1923	C5-3	95.8×96.6	18.40	Yes	1012750059
1961	880 Third Ave	1964	C6-6	100×95	14.78	Yes	1013080033
1961	1120 Sixth Ave	1962	C6-6	200.8×240	12.76	Yes	1012590001
1961	845 Third Ave	1964	C6-6	200.8×110	15.23	No	1013250001
1961	675 Third Ave	1966	C5-3	200.4×105	15.58	No	1013160001
1961	280 Park Ave	1962	C5-3	200.8×124.3	15.59	No	1012840033
1982	425 Lexington Ave	1987	C5-2.5	200.8×195	17.86	No	1012980023
1982	43 E 53 St	1990	C5-3	80×100.4	14.06	Yes	1012890028
1982	510 Madison Ave	2007	C5-3	100×157.5	18.58	Yes	1012880056
1982	505 Fifth Ave	2004	C5-3	100.5×133	18.80	Yes	1012770002
1982	563 Fifth Ave	1991	C5-3	100.4×180	19.09	Yes	1012820001

Table B.1 — The retrospective cohort: fifteen as-of-right Midtown buildings, five per regime, with tax-lot identifiers. Compl. is completion year and ZD is zoning district; Lot (ft) is frontage by depth, except the Chanin Building's three frontages. BBL is the borough–block–lot key, the city's unique tax-lot identifier, as recorded in MAPPLUTO. The Chanin Building is a corner-through lot and irregular: its 197.5 ft by 175 ft envelope would enclose 34,563 sq ft against a recorded lot area of 29,625, and at a gross floor area of 862,071 sq ft it carries the highest ratio in the cohort, 29.10: setback massing under a regime with no floor-area ceiling, not bonused bulk. Frontage scores may exceed 100% where the below-70° credit outruns blockage, by the construction of §81-274 (§3.9).

B.3 Setback verification (1916 cohort)

Approval under the 1916 regime cannot be read from a completion date alone: a building finished in the 1916-to-1961 window might have been designed and permitted before the resolution and grandfathered or altered since. For the five 1916 buildings the setback form was therefore confirmed directly, by inspecting each building's mass in the three-dimensional model and checking that it steps back above the street wall in the manner the sky exposure plane requires, rather than inferred from the year of completion. The Chanin Building (the pilot) and 364 Lexington Avenue record the explicit checks, the latter's stepped form verified in the Rhino model; regulatory approval under the 1916 regime, a fact of permit history the massing cannot show, is assumed for the cohort, all of it completed by 1960, inside the screen window of B.1. Where a candidate's setback form could not be confirmed, the systematic screen of B.1 supplies alternates from the same era and district.

B.4 Regime cohort sizes

The screen of B.1 also returns the size of each regime's as-of-right office stock in Midtown, the data behind the regime findings of §4.4. The counts are markedly uneven (Table B.2).

Measure	1916 setback	1961 tower-on-base	1982 DEC
Total office lots (n)	306	47	50
Mid-rise, ≥ 15 floors	214	41	40
As-of-right base-envelope tier	95*	5	17
Bonused / assemblage tier	—	36	23
Median bonused FAR	—	21.4	27.1

Table B.2 — As-of-right office stock in Midtown by regime (community districts 104–106), from the MAPPLUTO screen of B.1.

** Regular-lot count: screened 1916 lots of fifteen storeys or more carrying MAPPLUTO's regular-lot flag (IrrLotCode = N). The 1916 regime had no floor-area ceiling, floor area arrived with the 1961 resolution, so its tier is characterised by lot size (median 13,237 sq ft) and height (median 21 floors) rather than by an as-of-right FAR band; its floor-area ratios are an artefact of setback massing, not a regulatory target.*

B.5 Frontage parameters and the vantage construction

Each lot edge is tagged as street or lot line and carries its documented street width, from which the scorer derives the vantage half-width and selects the profile curve. Widths are taken from the city grid record and not from the model: the numbered avenues (Third, Fifth, Sixth) are 100 ft; Park Avenue is 140 ft; Lexington Avenue and Madison Avenue south of 42nd Street are 75 ft, Madison north of 42nd 80 ft; 42nd and 57th Streets are 100 ft major crosstowns; the remaining cross-streets are 60 ft. Streets narrower than 60 ft are evaluated on the 60 ft chart, the narrowest §81-272 provides, at their true half-width, a documented convention of the implementation. An intake record had invoked it, carrying E 53rd Street at 510 Madison Avenue as 50 ft; the tax-map lot-line gap measures approximately 54 ft, consistent with the 60 ft grid width less digitization tolerance, and the frontage is scored at 60 ft, matching the facing frontage of 43 E 53rd Street on the same segment. Width is the single most consequential input: relabelling one Madison frontage from 100 ft to 80 ft moves that building's score by more than twenty points.

Two constructions in the ZR depart from the geometric centreline, and both bear on this cohort. On a street 75 ft wide, the centreline is fixed at 40 ft from the front lot line, and on a street more than 100 ft wide at 50 ft. The implementation of §3.9 places the vantage

at half the fed width. Park Avenue was therefore entered as 100 ft, which reproduces the ZR construction exactly; the three 75 ft frontages in the 1916 cohort (Chanin, 364 Lexington, 292 Madison) were entered at their true width and so evaluated at 37.5 ft rather than 40 ft, which depresses those frontage scores. Those three frontages were therefore re-scored with 80 ft fed in place of 75, which reproduces the regulatory vantage of 40 ft on the same 75/80 ft chart and leaves the chart selection untouched. The correction raises the Chanin Building by 1.0 points, 364 Lexington Avenue by 1.6, and 292 Madison Avenue by 3.1, carrying the last from 66.1 to 69.2 and so across the floor of the acceptance band. The Lexington Avenue frontage of 425 Lexington Avenue, admitted after the correction, was entered at 80 ft from the outset. The evolutionary matrix retains the geometric half-width, uniformly across regimes; §5.3 records the consequence.

B.6 Cohort scores

Table B.3 records the score returned for each building by the same compiled components used in the evolutionary evaluation. Every aggregation was recomputed from the per-frontage record of Table B.4 and agrees with the scorer to the last reported digit.

Regime	Building	DEC (%)	Min street (%)	Mean SVF	Legal, 1982
1916	Chanin Building	58.9	41.0	0.730	No
1916	364 Lexington Ave	62.3	53.1	0.758	No
1916	2 Park Ave	49.2	37.0	0.704	No
1916	485 Madison Ave	44.9	40.1	0.747	No
1916	292 Madison Ave	69.2	64.0	0.783	No
1961	880 Third Ave	97.7	95.1	0.816	Yes
1961	1120 Sixth Ave	89.5	86.4	0.733	Yes
1961	845 Third Ave	89.0	78.7	0.767	Yes
1961	675 Third Ave	85.3	83.1	0.823	Yes
1961	280 Park Ave	76.0	68.4	0.769	Yes
1982	425 Lexington Ave	78.6	73.0	0.769	Yes
1982	43 E 53 St	78.1	78.1	0.760	Yes
1982	510 Madison Ave	79.1	75.4	0.770	Yes
1982	505 Fifth Ave	79.9	76.9	0.808	Yes
1982	563 Fifth Ave	82.0	78.8	0.764	Yes

Table B.3 — Retrospective cohort scores. Legality is the conjunction of the 75% overall and 66% per-street tests of §81-274(i). 505 and 563 Fifth Avenue are scored against the unmodified evaluation, which is not the rule their zoning lots were held to (B.8).

B.7 Per-frontage record

Table B.4 gives the frontage-level record on which Table B.3 and the built-record correlation of §4.3 rest. The SVF column is the mean over the sample points on that frontage's centreline.

Building	Frontage	Width (ft)	Chart	DEC (%)	Mean SVF	n
Chanin Building	E 42nd St	100	100	84.7	0.800	13
Chanin Building	Lexington Ave	80	80	58.5	0.716	22
Chanin Building	E 41st St	60	60	41.0	0.697	19
485 Madison Ave	Madison Ave	80	80	50.9	0.775	11
485 Madison Ave	E 52nd St	60	60	40.1	0.724	13
364 Lexington Ave	E 41st St	60	60	53.1	0.716	16
364 Lexington Ave	Lexington Ave	80	80	76.3	0.821	11
2 Park Ave	E 33rd St	60	60	37.0	0.685	22
2 Park Ave	Park Ave	100	100	55.7	0.733	22
2 Park Ave	E 32nd St	60	60	54.8	0.694	22
292 Madison Ave	E 41st St	60	60	74.4	0.775	10
292 Madison Ave	Madison Ave	80	80	64.0	0.791	11
880 Third Ave	Third Ave	100	100	100.2	0.844	11
880 Third Ave	E 53rd St	60	60	95.1	0.787	10
1120 Sixth Ave	W 43rd St	60	60	87.7	0.714	25
1120 Sixth Ave	Sixth Ave	100	100	95.2	0.778	21
1120 Sixth Ave	W 44th St	60	60	86.4	0.713	25
845 Third Ave	E 51st St	60	60	80.3	0.760	12
845 Third Ave	Third Ave	100	100	98.8	0.772	21
845 Third Ave	E 52nd St	60	60	78.7	0.764	11
675 Third Ave	E 42nd St	100	100	85.6	0.882	9
675 Third Ave	Third Ave	100	100	83.1	0.791	21
675 Third Ave	E 43rd St	60	60	89.1	0.840	11
280 Park Ave	E 49th St	60	60	83.2	0.782	12
280 Park Ave	Park Ave	100	100	68.4	0.757	21
280 Park Ave	E 48th St	60	60	82.7	0.779	12
425 Lexington Ave	E 43rd St	60	60	73.0	0.773	20
425 Lexington Ave	Lexington Ave	80	80	80.8	0.771	21
425 Lexington Ave	E 44th St	60	60	82.7	0.760	16
43 E 53 St	E 53rd St	60	60	78.1	0.760	9
510 Madison Ave	E 53rd St	60	60	75.4	0.745	17
510 Madison Ave	Madison Ave	80	80	84.7	0.810	11
505 Fifth Ave	E 42nd St	100	100	76.9	0.792	15
505 Fifth Ave	Fifth Ave	100	100	84.2	0.831	11
563 Fifth Ave	E 46th St	60	60	78.8	0.725	19
563 Fifth Ave	Fifth Ave	100	100	87.6	0.832	11

Table B.4 — Per-frontage scores and sky measurements, thirty-six frontages across the fifteen buildings. The two collinear segments of the Third Avenue frontage of 675 Third Avenue are merged, as one street.

B.8 Compliance routes and subdistricts

A building inside the Special Midtown District complies with the height and setback rules by one of two elective routes: §81-26, daylight compensation, or §81-27, the DEC. The election is recorded on the approved zoning diagram in the Department of Buildings job folder rather than in any public dataset. Only a building that elected §81-27 was held to the score this thesis computes. The routes were not recoverable for the 1982 cohort from open sources; for 563 Fifth Avenue the drawings predate digitisation. The audit therefore reports the score each building delivers rather than the test it passed. The 1916 and 1961 buildings predate the DEC entirely, so for them the question does not arise: their scores are counterfactual by construction, which is what the calibration audit of H1 requires.

Two lots are subdistrict-modified. Under §81-253 the provisions of §§81-26 and 81-27 are supplemented and modified by §§81-81 and 81-83 in the Fifth Avenue Subdistrict, and by §§81-71 and 81-75 in the Theater Subdistrict. Map 1 of Appendix A to the ZR places 505 and 563 Fifth Avenue within the Fifth Avenue Subdistrict, whose shaded band runs along both sides of Fifth Avenue to a depth of 150 to 200 ft. Neither was therefore held to the evaluation as this thesis implements it. Both are retained: §81-83 modifies the street wall rather than the scoring, and both clear the thresholds. The modification is recorded, not corrected for.

Carnegie Hall Tower, originally scored as part of this cohort, was removed at a late stage because it fails the first gate of §3.8. It stands on a city-owned lot and required approval from the Landmarks Preservation Commission, Community Board 5, the City Planning Commission and the Board of Estimate, which granted it on 17 September 1987; its development rights were transferred from the adjoining Carnegie Hall landmark lot through a merged zoning lot. It also lies within the Theater Subdistrict, where §81-253 modifies §§81-26 and 81-27, a second, independent ground on which it was never held to the unmodified chart. It is one of the discretionary cases the gate was written to exclude.

B.9 Screening annex: the 1982 replacement

When Carnegie Hall Tower failed the as-of-right gate, the rule for replacing it was fixed in writing before any candidate was scored, so that the choice could not be made on the strength of a Daylight Evaluation score. The gates were: approval under the 1982 regime; no discretionary approval and no development rights transferred by merger; a built FAR no more than 1.30 times the district maximum recorded in MAPPLUTO; no

street frontage over 250 ft; and at least eighteen storeys. Among candidates passing every gate, the one whose lot typology was absent from the 1982 cohort entered first, ties being broken on earliest completion. Membership of a subdistrict was recorded as a flag, not a gate, consistent with the treatment of 505 and 563 Fifth Avenue.

The floor-area screen is the operative one, and it needs no estimation: MAPPLUTO reports both the built ratio and the district maximum for commercial use, so the quotient is read directly. The four buildings retained from the original cohort return 1.17, 1.24, 1.25 and 1.27; the threshold of 1.30 was set from them. Table B.5 records every candidate screened.

The screen is itself a result. Of five 1982 candidates, three carried a floor-area ratio between 1.42 and 2.19 times their district maximum, and two were approved before the chart existed. The tower the 1982 regime actually produced in Midtown was assembled, bonused, or both: the same bimodal structure §4.4 finds in the 1961 stock, arriving independently from the regime that replaced it. That Carnegie Hall Tower scored 61.4% and Park Avenue Tower 90.7%, and that neither figure entered the decision, is the substance of the protocol rather than a coincidence of it.

Candidate	Year	Built FAR	Max FAR	Ratio	DEC	Determination
425 Lexington Ave	**1987**	**17.86**	**15**	**1.19**	**78.6**	**Admitted: corner-through, the absent typology**
Tower 49	1985	26.25	12	2.19	89.0	Excluded: merged zoning lot
599 Lexington Ave	1985	21.87	15	1.46	49.9	Excluded: transit-improvement special permit
535 Madison Ave	1981	26.20	15	1.75	−8.1	Excluded: predates 1982 regime, assembled lot
Park Ave Tower	1986	17.07	12	1.42	90.7	Excluded: floor area exceeds the C5-2.5 maximum
550 Madison Ave	1984	18.62	15	1.24	—	Excluded: designed in 1978, approved before the regime

Table B.5 — The 1982 replacement screen. Ratio is Built FAR divided by the maximum commercial FAR recorded in MAPPLUTO. DEC scores are given for the five candidates modelled; they played no part in any determination and are reported here so that the screen cannot be mistaken for a choice among outcomes. 550 Madison Avenue, the AT&T Building, was excluded on its design date before it was modelled. A negative score arises where blockage and penalty exceed available daylight; the scorer does not floor at zero.

Appendix C: Instrument Validation

The two measuring instruments the study relies on, the DEC scorer and the HV calculator, are each validated against a target with a known answer before they are trusted. This appendix reports the scorer's reproduction of the regulation's own worked example, decomposed to its regulatory components; the specification of the sky-view-factor sampling that supplies the independent openness measure; and the two gates that verify the HV calculation.

C.1 DEC scorer validation against the worked example

The scorer was run on the worked example published in the ZR and reproduced by Kwartler and Masters (1984, Fig. 10, p. 182): a one-hundred-foot street with a lot one hundred and seventy feet wide and one hundred feet deep, carrying three slabs across its full frontage: a twenty-foot-deep street wall one hundred and forty feet tall, a tower from twenty to eighty feet deep rising to three hundred and twenty feet, and a twenty-foot-deep rear element twenty-three feet tall. The published score sheet gives available daylight 89.9, blockage above seventy degrees −20.5, credit below seventy degrees +0.0, and a profile penalty of −0.4 (0.1 × 1.5 + 0.25 × 1.0), for an overall score of 76.75% (the current ZR text prints −0.45 and 76.70; the 1984 print is used as the anchor throughout). The implementation returns available 90.0, blockage −23.7, below-seventy +0.0, profile −0.70, and 72.89% (Table 3.1). Each component reconciles. The available-daylight count matches to a tenth of a square, the residual a partial-column counting convention at the boundary. The below-seventy credit matches exactly, though the term is zero in the published sheet and in the implementation alike, so the worked example does not exercise it.

The blockage term is reconciled subsquare by subsquare. The canonical rule adopted here is the strict any-portion rule of §81-274(a)(3), charging the entire subsquare wherever a building blocks any portion of it, which is also the rule of the production algorithm published by Fishhaut, Weingarten and Neiger (1984, pp. 185–186). It reads 23.7 squares against the published 20.5, the excess falling on the grazing rows that the 1984 hand drafting tolerated. The alternative centre-point rule, which charges a subsquare only when its centre lies within the projected silhouette, reads 21.8 squares and yields 75.0%; it has no published anchor and is retained solely as the sensitivity variant of C.9. The profile penalty of −0.70 is computed against the width-specific nineteen-point digitisation of the official chart (C.4), the published −0.4 reflecting the drawn curve's position on the hand-drafted sheet. The two conventions carrying the

gap, namely the blockage rule and the profile digitisation, are both conservative, and the aggregate never reads a higher score than the 1984 published practice, so a massing the scorer certifies compliant would also have passed the drafting the regulation was calibrated against. That one-directional bias is what licenses the acceptance band of §3.8, [68, 76], where cohort scores are read raw against this worked example as an anchor, with no correction applied.

One consequence deserves recording. Under the canonical rule the regulation's own worked-example building scores 72.89% and therefore fails the 75% overall threshold it was published to demonstrate, while still passing the 66% per-street test. The convention band straddles the legal pass mark on the very building the ZR uses to teach the method. No result in this thesis depends on which side of that mark the example falls (C.9), but the fact is the sharpest available demonstration that a daylight evaluation is not fully specified by its published arithmetic, and that the counting convention must be declared before any score can be compared with any other.

C.2 SVF sampling specification

The SVF is computed by the CenterlineSVF component (Appendix E) as the openness measure independent of the Chart. At each sample the upward hemisphere is discretised on an altitude-by-azimuth grid of rays, and the SVF is the fraction of that hemisphere left unobstructed by the building. Samples are placed along each street's centreline at a pedestrian eye height of 1.5 m and a spacing of 3.0 m, with at least five samples on every frontage so that short frontages are not under-sampled; the street value is the mean over its samples, and the building's value the frontage-weighted mean over its streets. The rays test the evolved or modelled building alone, with no surrounding context mass; this is the single-building convention the DEC adopts, since the Chart too plots only the building under evaluation, and it ensures that any divergence between the two measures is a difference in how each instrument samples the sky, not in what each is permitted to see. The measure is never an objective and no form is ever selected for it; its role is diagnostic, as set out in §§3.5 and 3.6.

Ray-grid resolution policy. The grid above was originally 72 azimuths by 18 elevation bands (1,296 rays). A pre-registered convergence test against a doubled grid failed: per-sample values moved by up to 0.013 and building-level means by 0.002, with error oscillatory in sign at the sample level. A resolution ladder was run against two analytic anchors: the open scene, exact at any resolution by self-normalisation, and the infinite canyon at two depths against Oke's (1981) closed form. It showed 288 by 72 (20,736

rays) converged to within 5×10^{-4} at the mean, confirmed independently by a fine-grid reference. The retrospective evaluation is measured at 288 by 72. The evolutionary evaluation retains 72 by 18: a full re-measurement of the 1982 corner cell at 288 by 72 moved its 208 per-individual means by at most 0.0023, its off-front rank correlation from +0.868 to +0.864, its on-front correlation by 0.0001, and its legal-minimum sky value by 0.0005. No printed conclusion depends on shifts of this size, and rank statistics tolerate them by two orders of magnitude. Per-sample profile values carry ±0.01 resolution uncertainty at 72 by 18 and ±0.004 at 288 by 72; along-frontage profiles reported in this thesis use the 288-by-72 build.

C.3 HV computation and validation

The HV calculator (Appendix E) measures the size of each regime's evolved design space as the two-dimensional area of objective space dominated by its Pareto front relative to a fixed reference point. Both objectives are minimised, as the negated products of daylight score and legality, and of volume and legality, the 1982 regime's daylight objective alone carrying no multiplier (§3.7). The reference point is the origin, which every feasible individual dominates; the same reference is used for all sixteen cells so that the hypervolumes are comparable, the daylight coordinate is clipped at zero, and any point at or beyond the reference is excluded. The calculation is checked by two gates before use. The first runs a known point set whose dominated area is 5,500,000 computed by hand: the implementation must return exactly that. The second appends a point the front already dominates, (−70, −40000), and requires the HV to be unchanged and the non-dominated count to remain three. This confirms that dominated points neither inflate the measure nor disturb the front.

C.4 The published Daylight Evaluation Diagrams

The three figures below reproduce the DEC as published by the NYC DCP, one chart per charted street width. Each is the curvilinear projection of the sky dome from the regulatory vantage that §81-272 defines. These are the instruments the rectilinear scorer of Appendix C.1 reproduces; the 19-point profile curves it digitises (PROFILE_CURVE_60 / _80 / _100) are traced from exactly these graphics.

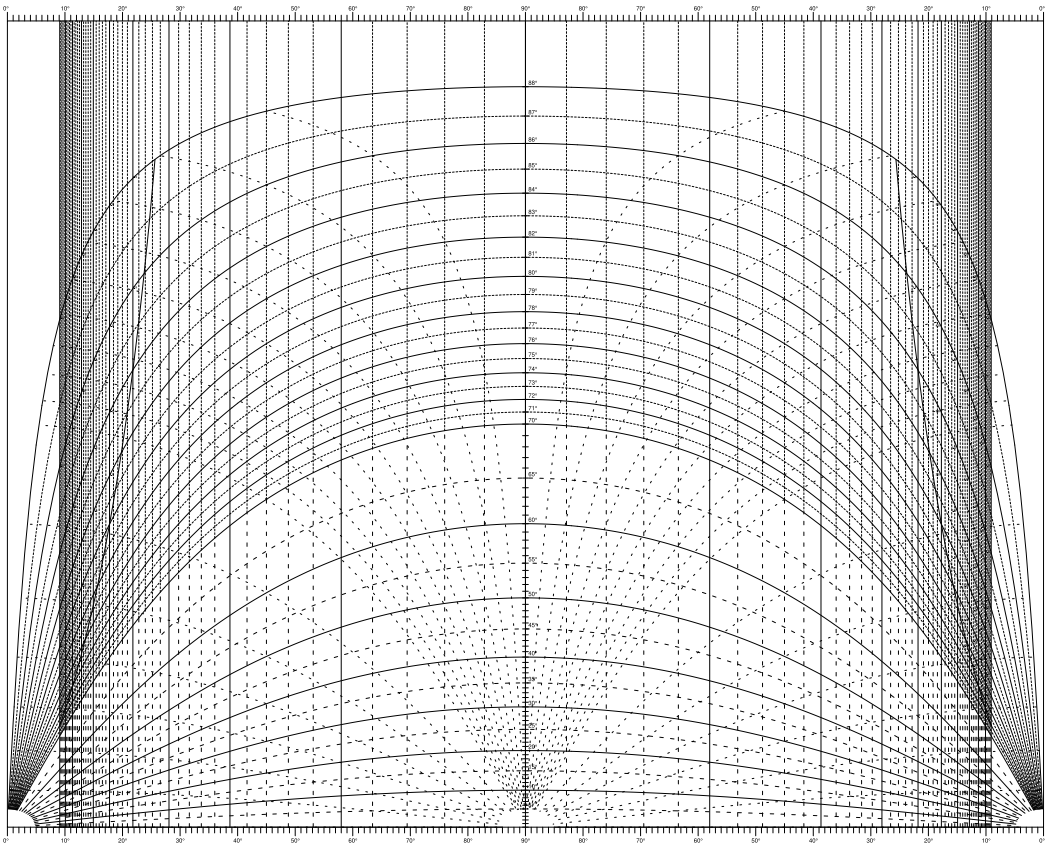


Figure C.1 — Daylight Evaluation Diagram: 60-foot street (NYC Department of City Planning).

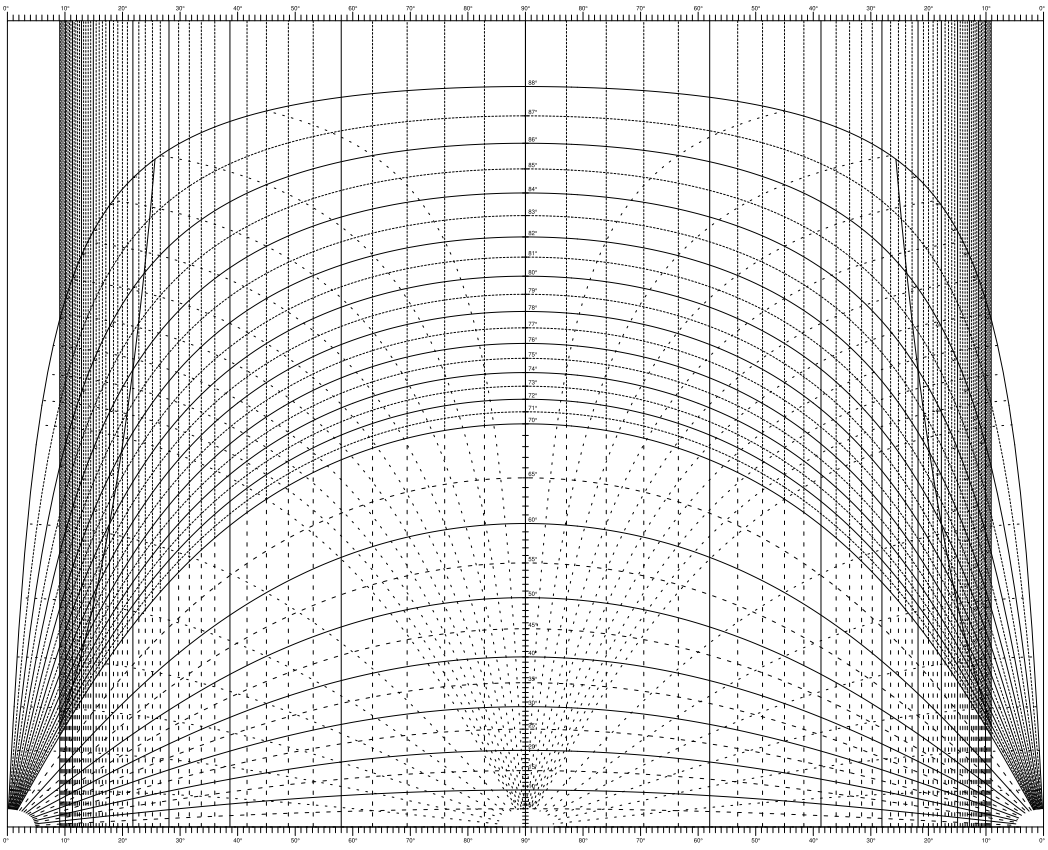


Figure C.2 — Daylight Evaluation Diagram: 80-foot street (NYC Department of City Planning).

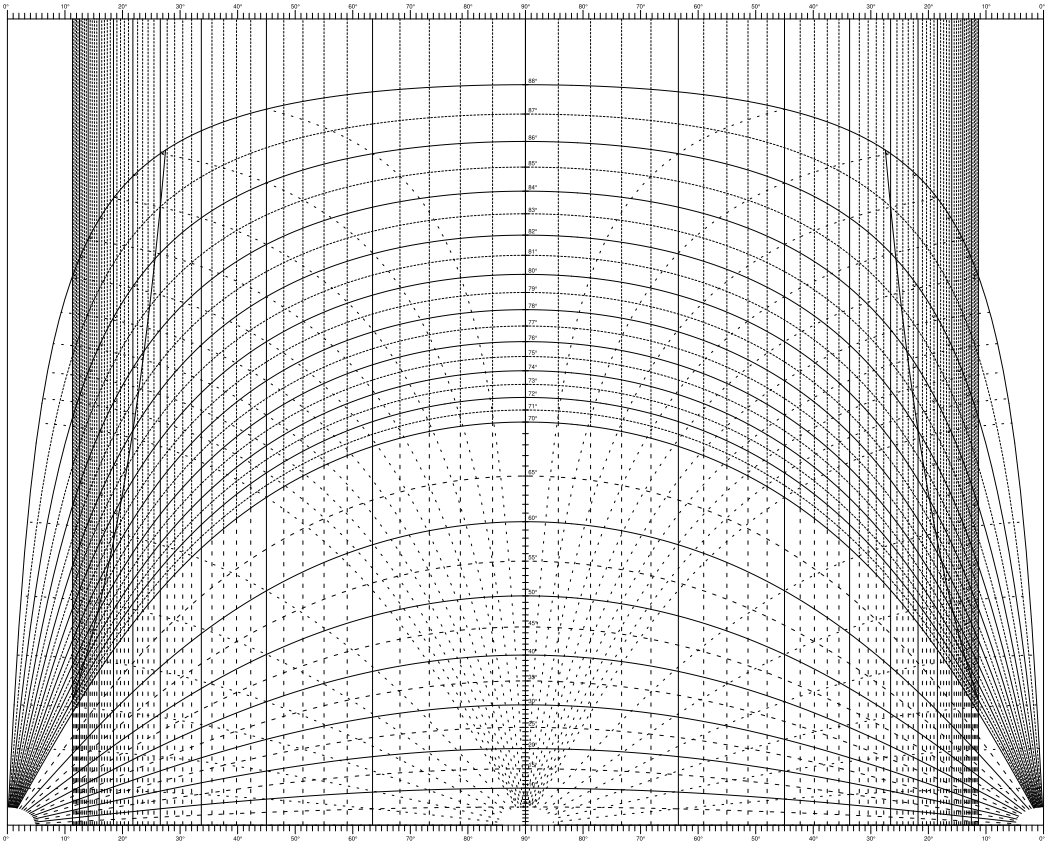


Figure C.3 — Daylight Evaluation Diagram: 100-foot street (NYC Department of City Planning).

C.5 The scorer algorithm

The scorer implements the rectilinear chart of Fishhaut, Weingarten and Neiger (1984). For each street frontage of the lot, two vantage points are constructed on the street centreline, each 250 feet from the intersection of the centreline with a far side lot line, per the plotting rules of §81-273, side lot lines being taken as the lot's non-parallel boundary edges, whether or not those edges are themselves street lines, the construction a corner lot requires. Every massing vertex is projected onto the chart plane by the rectilinear mapping, under which the picture-plane coordinates are the vertex's along-street distance and height scaled by the ratio of the half street width to the vertex's setback from the centreline, a mapping that sends straight edges to straight lines and reduces blockage testing to polygon–rectangle intersection. The chart surface above seventy degrees of elevation is rasterised into subsquares of five feet of frontage by one degree of elevation; a subsquare is counted blocked when any part of a projected face covers it, whether a vertex inside the cell, the cell centre inside a face, or a face edge crossing the cell; this is the strict any-portion rule of §81-274(a)(3) and of the SOM production algorithm (C.1). One property of the closed-interval semantics is documented from the verification programme of §3.9: where a silhouette edge falls within a microfoot of a

subsquare boundary, single-precision mesh storage decides which side it lands, and the charging of the adjacent column with it, a tie that is measure-zero for measured and evolved inputs, demonstrated and bounded by the symmetry-rig pair. The centre-point variant, which charges the cell only when its centre lies inside a face, is retained solely for the convention band of C.9. Below seventy degrees, unblocked full squares earn the credit the scoring rules prescribe, and the chart is trimmed at the near boundary line derived from the hundred-foot boundary depth of the standard evaluation setup.

C.6 The profile curve and the aggregation of scores

The profile penalty of §81-274 is triggered by the profile curve, the drawn line on the official chart graphic that rises from the intersection of the seventy-two-degree row with the far-lot-line vertical and represents the predominant built profile of Midtown streets. The published text does not give the curve in closed form; the implementation therefore encodes it as a digitised coordinate table with linear interpolation and applies the penalty weights of the regulation's table only to blocked subsquares whose far-lot-line edge lies beyond the curve at their elevation. The curve is a width-specific digitisation of the official chart graphics, a nineteen-point coordinate table for each charted street width of sixty, eighty and one hundred feet, selected on the frontage's width, traced from the published diagrams reproduced in Appendix C.4 and validated against the worked example. Scores from the two vantage points of each street are averaged, street scores are combined by frontage-weighted average into the overall score, and the per-street minimum is carried separately for the 66% test.

C.7 Component decomposition of the worked example

The validation decomposes the score into its regulatory components (Table 3.1). The available-daylight term matches the published value to within a tenth of a square, and the below-seventy-degree credit matches exactly, that term being zero in both. The blockage term reads 23.7 squares against the published 20.5 under the canonical any-portion rule of §81-274(a)(3), which charges a subsquare wherever the silhouette covers any part of it, and which is the rule of SOM's published production algorithm; the centre-point variant, which charges only when the subsquare centre lies within the silhouette, reads 21.8 and is retained solely for the convention band of C.9. The profile penalty reads 0.70 against the published 0.4, computed against the width-specific nineteen-point digitisation of the official chart. The aggregate score is 72.89% against the published 76.75, and the gap is conservative in aggregate: the implementation never reads higher than the published practice; C.9 gives the decomposition. For the evolutionary

evaluation this means no massing is certified compliant that the 1984 practice would have failed; for the retrospective evaluation it fixes the acceptance band of §3.8.

C.8 Sampling protocol for the coupling test

The correlation across the evolved populations is measured on a drawn sample rather than the whole population, constructed per cell to cover the daylight axis evenly. From each population the distinct genomes are taken once, since identical gene vectors decode to identical massing, and so to an identical SVF, and carry no further information; the individuals whose genomes lie on the Pareto front are set aside as a separate on-front sample. The remaining unique, off-front individuals are stratified by DEC score into fifteen equal-width bins, and up to eighteen are drawn from each bin at even score spacing; the resulting sample is ordered by score and paired positionally with the SVF returned for it, an alignment checked per cell against a known individual. Stratifying by score rather than sampling at random keeps the dense middle of the distribution from dominating the rank correlation, so that the coefficient describes the relationship across the whole score range; it cannot manufacture a correlation, since even coverage of an unrelated axis would still return a coefficient near zero. The off-front set serves as the primary evolutionary test precisely because these are the individuals the algorithm did not select onto the front, so their score-to-openness correlation is the least exposed to the circularity described above; the on-front sample is reported alongside it.

C.9 Convention-band sensitivity of the verdicts

The 3.86-point gap of C.1 is not distributed evenly across the score's four terms. On a common denominator the blockage rule carries 3.56 points of the gap and the digitisation of the profile curve 0.33; the below-seventy credit contributes nothing, that term being zero on both sheets, and the available-daylight count works 0.03 points in the opposite direction, the implementation reading marginally more available sky than the hand-drafted sheet. Since both are choices rather than measurements, the question a reader is entitled to ask is not whether the implementation reproduces one number but whether any conclusion in this thesis depends on the choices made.

The two blockage rules bound the answer. The canonical any-portion rule reads the worked example at 72.89%; the centre-point rule, which charges a subsquare only when its centre lies within the silhouette, reads it at 75.0%. Because any-portion charges every subsquare centre-point charges and the grazing ones besides, the canonical reading is the lower of the two by construction. Applying the 2.11-point band uniformly to the

retrospective cohort tests every verdict against the width of the convention, and Table C.1 reports the outcome. The calibration verdict of §4.1 is invariant: four of the five 1916 buildings lie below the acceptance band under either rule and the fifth, 292 Madison Avenue, sits just inside its floor under both; the 1961 cohort lies at or above the band; and the 1982 cohort lies above it under both. The coupling verdict of §4.3 is invariant by construction: a uniform shift is a monotone transformation, and the Spearman coefficient is a rank statistic, so the built-record correlation stands at +0.565 under both readings, to three decimal places. Legality under the 75 and 66% thresholds is unchanged for all fifteen buildings.

One building crosses the band boundary, and it carries no hypothesis. 280 Park Avenue sits exactly at the band ceiling of 76.0 under the canonical rule and rises above it under the centre-point variant. Nothing else moves: 292 Madison Avenue is inside the band under both readings, and the remaining thirteen buildings stay on the side of the band they began on. The regime-level readings of §4.1 stand under either convention: the 1916 stock materially below the band the regulation implies for it, the 1961 stock materially above.

Regime	Canonical, any-portion	Sensitivity, centre-point	In band, canonical	In band, centre-point	Legal under 1982
1916	56.9	59.0	1/5	1/5	0/5
1961	87.5	89.6	1/5	0/5	5/5
1982	79.5	81.7	0/5	0/5	5/5

Table C.1 — Verdict sensitivity across the blockage-convention band. Means of the overall DEC score per regime; the centre-point column applies the 2.11-point bound of C.1 uniformly. The acceptance band is [68, 76] (§3.8). No hypothesis verdict changes: the one building that crosses the band boundary is named in the text.

Appendix D: Evolved Design Spaces in Full

This appendix holds the complete evolutionary evaluation evidence behind Chapter 4: the design-space envelopes overlaid by regime, the convergence and coupling diagnostics, the evolved population and Pareto front of every one of the sixteen cells organised by regime, and the full numerical results. All sixteen runs share identical settings, a population of one hundred evolved for one hundred generations from a single seed under NSGA-II, so that any difference between cells is attributable to the regime or the lot, never to the optimiser. Hypervolume is computed on the convention fixed in Appendix C: reference point at the origin, DEC score clipped at zero.

D.1 Design-space envelopes by regime

Each panel overlays the four regimes' Pareto fronts on one lot. The nested shaded areas are the hypervolumes Table 4.1 reports as ratios; the fronts themselves show where the envelopes differ in shape, not only in size. The dot on each front marks its left-extent: the lowest DEC score at which the regime still holds a non-dominated form. Only the unregulated baseline reaches deep into the low-daylight, high-bulk region; under the two prescriptive regimes the decoder's clipping means such forms never exist to be plotted, and under the 1982 chart the front stops at the 75 line, three of the four cells holding no sub-threshold form at all and the corner-through a single one, at 74.3 (§4.2). Regulation does not merely shrink the envelope; it removes the option of trading daylight for bulk, the envelopes by construction and the chart by score.

D.2 Convergence and coupling

Two diagnostics stand behind the results. The first confirms the fixed generation budget was sufficient: every cell reaches at least 99% of its final HV within the hundred generations, the slowest, namely the 1982 through-block, arriving at generation 94, and no regime systematically slower than the rest. The second summarises H3 across the whole matrix: the rank coupling between DEC score and centreline SVF sits above 0.85 in every cell off-front, the load-bearing sample, and on-front in fourteen of sixteen, the two soft cells noted at Table 4.3.

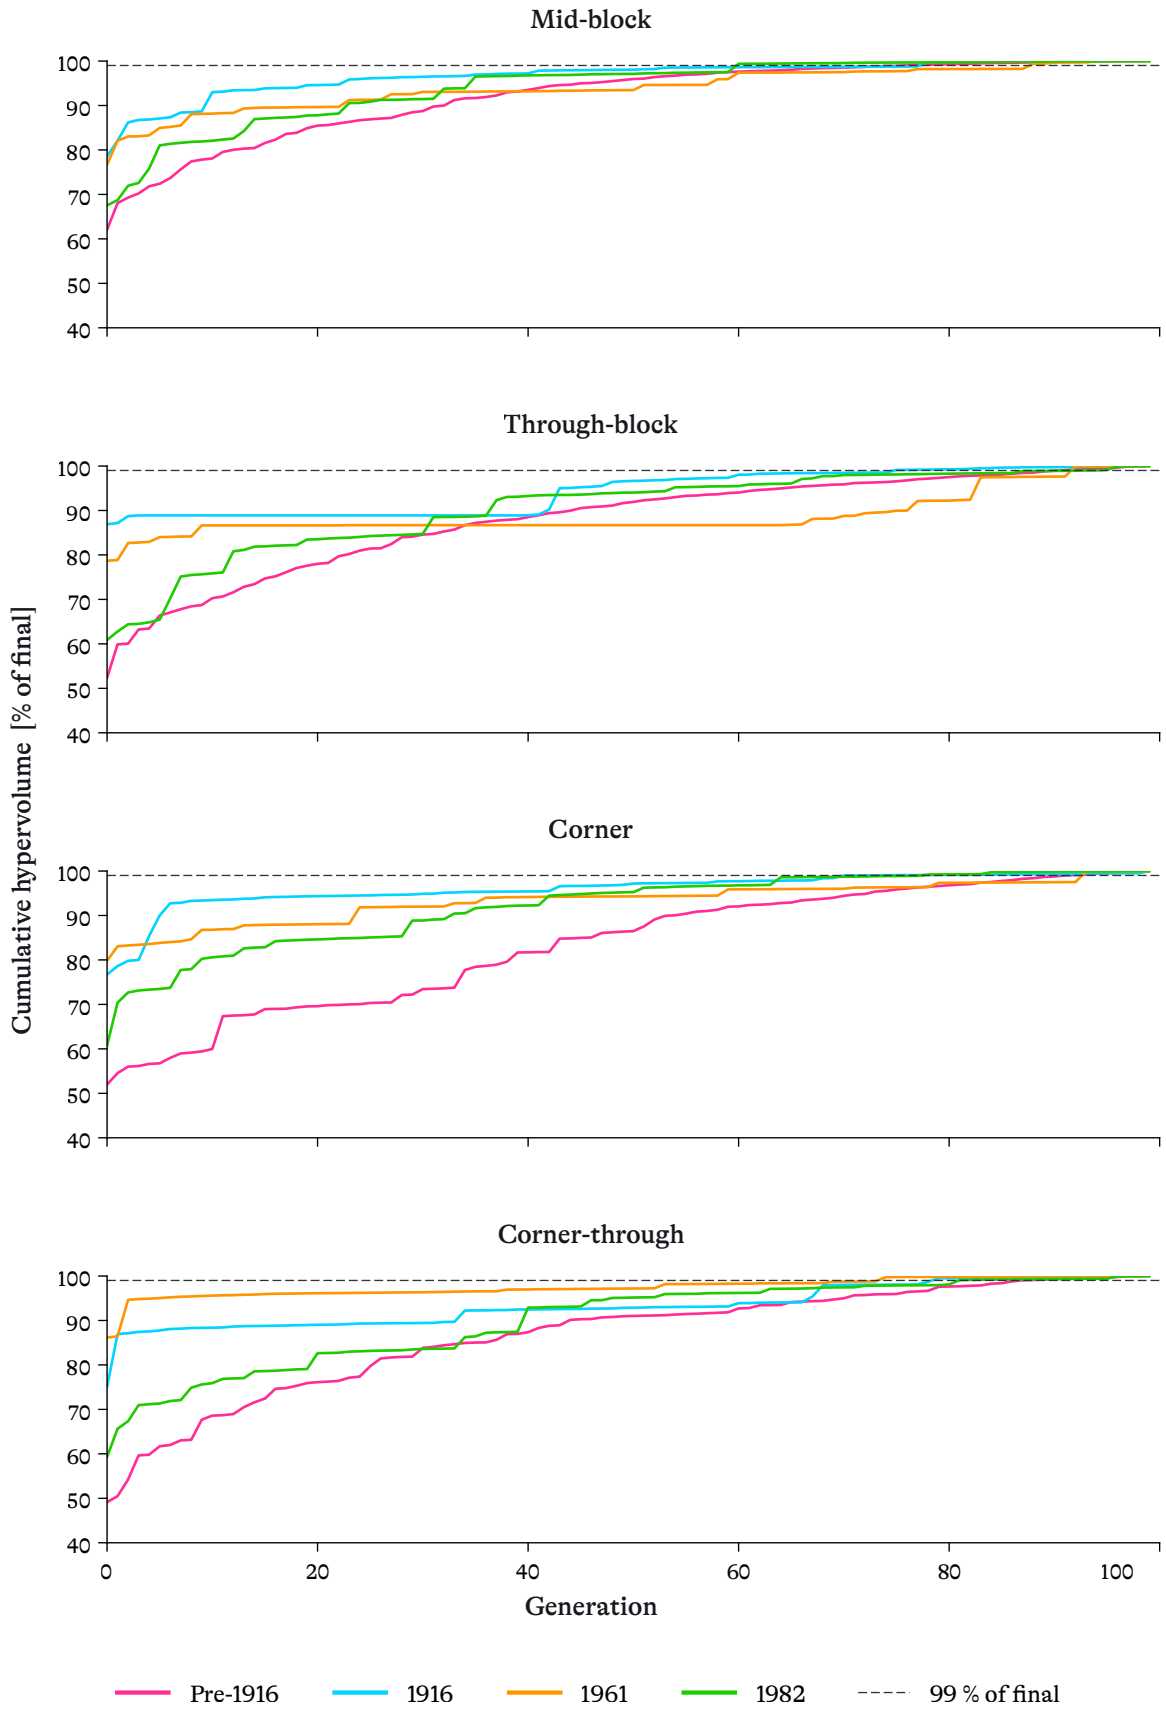


Figure D.1 — HV convergence by regime, overlaid within each typology; the dashed line marks 99% of final.

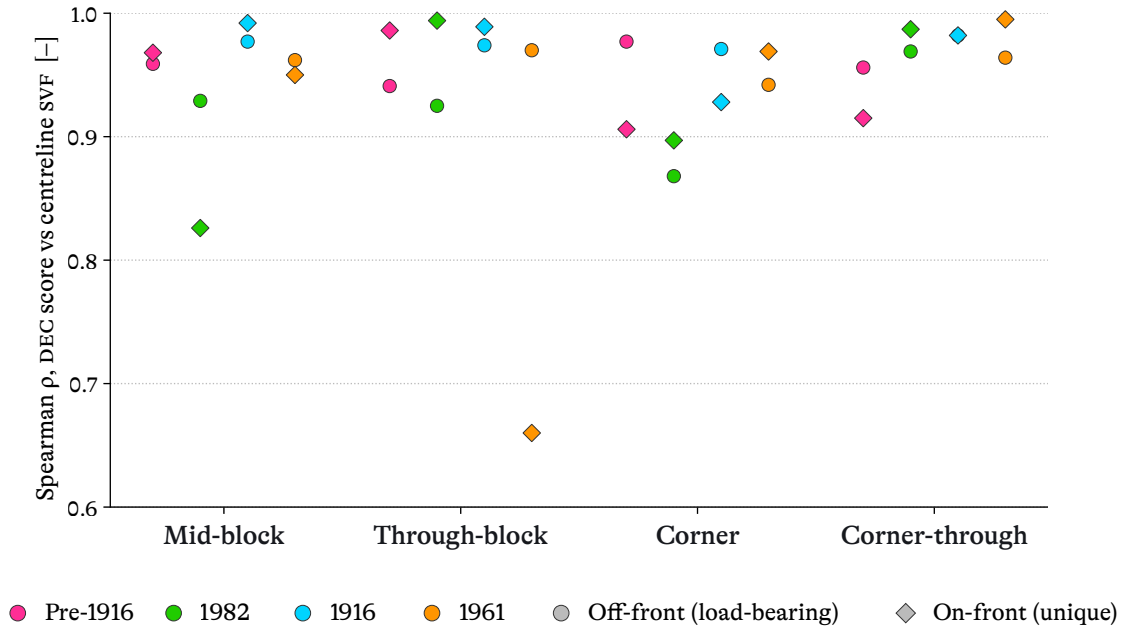


Figure D.2 — DEC–SVF rank correlation, all sixteen cells: Spearman rank correlation between DEC score and centreline SVF; filled circles off-front (the load-bearing test), diamonds on-front.

D.3 Design spaces by regime

The four figures below (Figures D.3–D.6) show, for each regime in turn, the evolved population (light points) and the recovered Pareto front on all four lot typologies. Read together they make the envelope differences concrete: the unregulated baseline populates the full daylight range, the 1982 chart tracks it closely at high daylight, and the two prescriptive regimes fall away, the 1961 envelope most sharply, its front barely leaving the high-score corner.

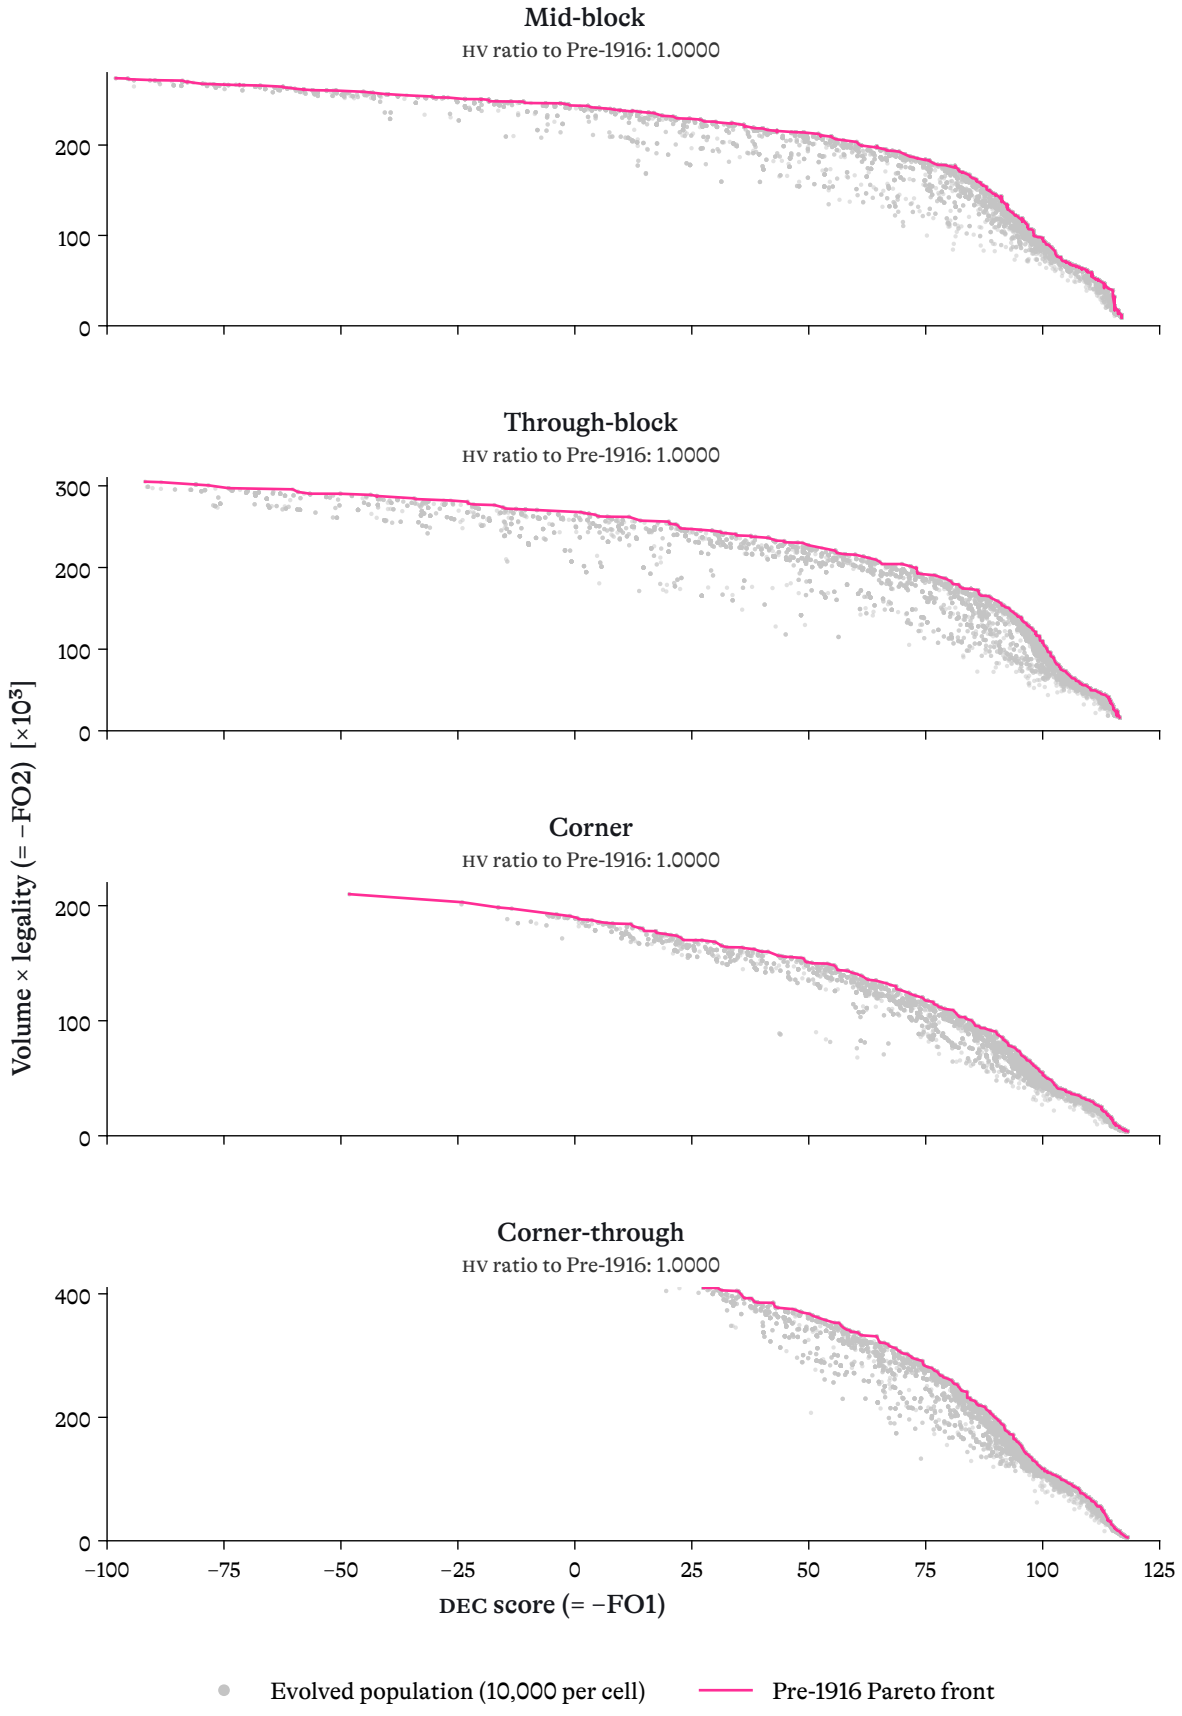


Figure D.3 — Pre-1916: evolved population (light) and Pareto front by typology, with the HV ratio to the pre-1916 baseline noted per panel.

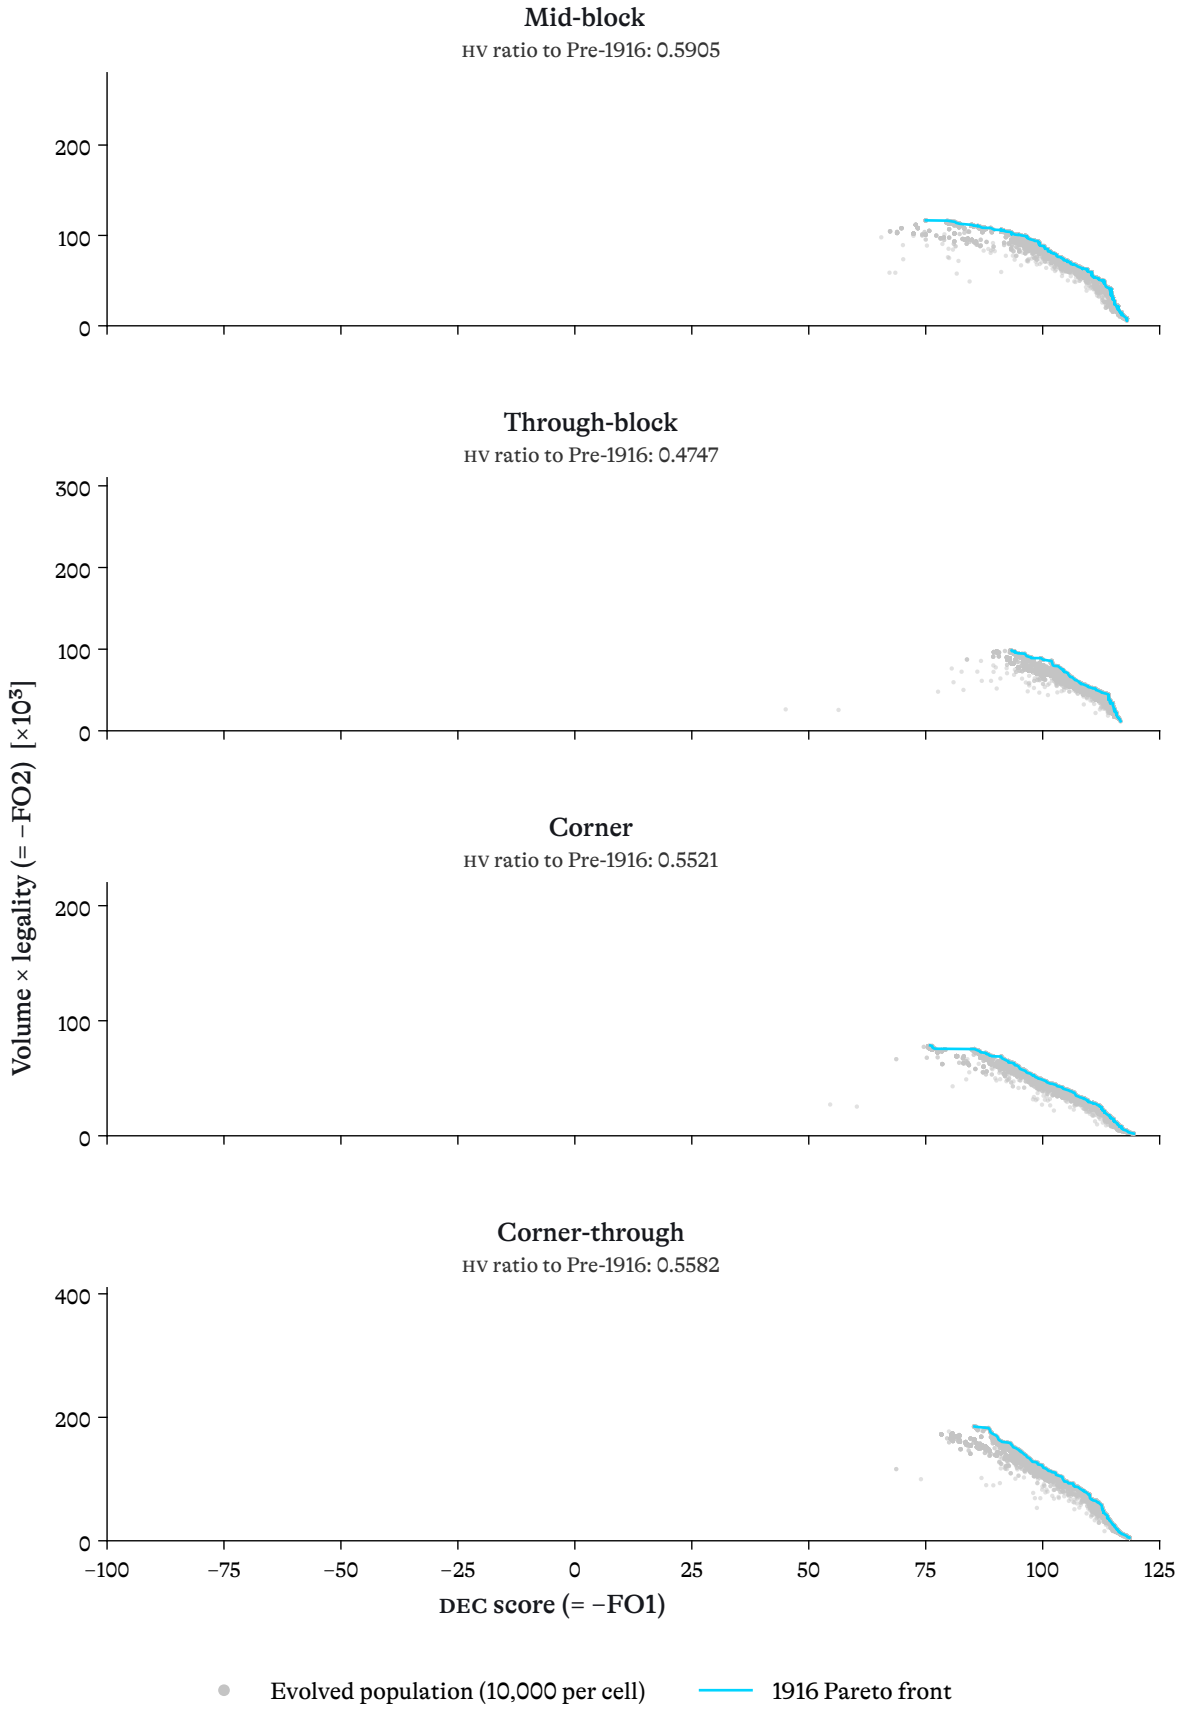


Figure D.4 — 1916: evolved population (light) and Pareto front by typology, with the HV ratio to the pre-1916 baseline noted per panel.

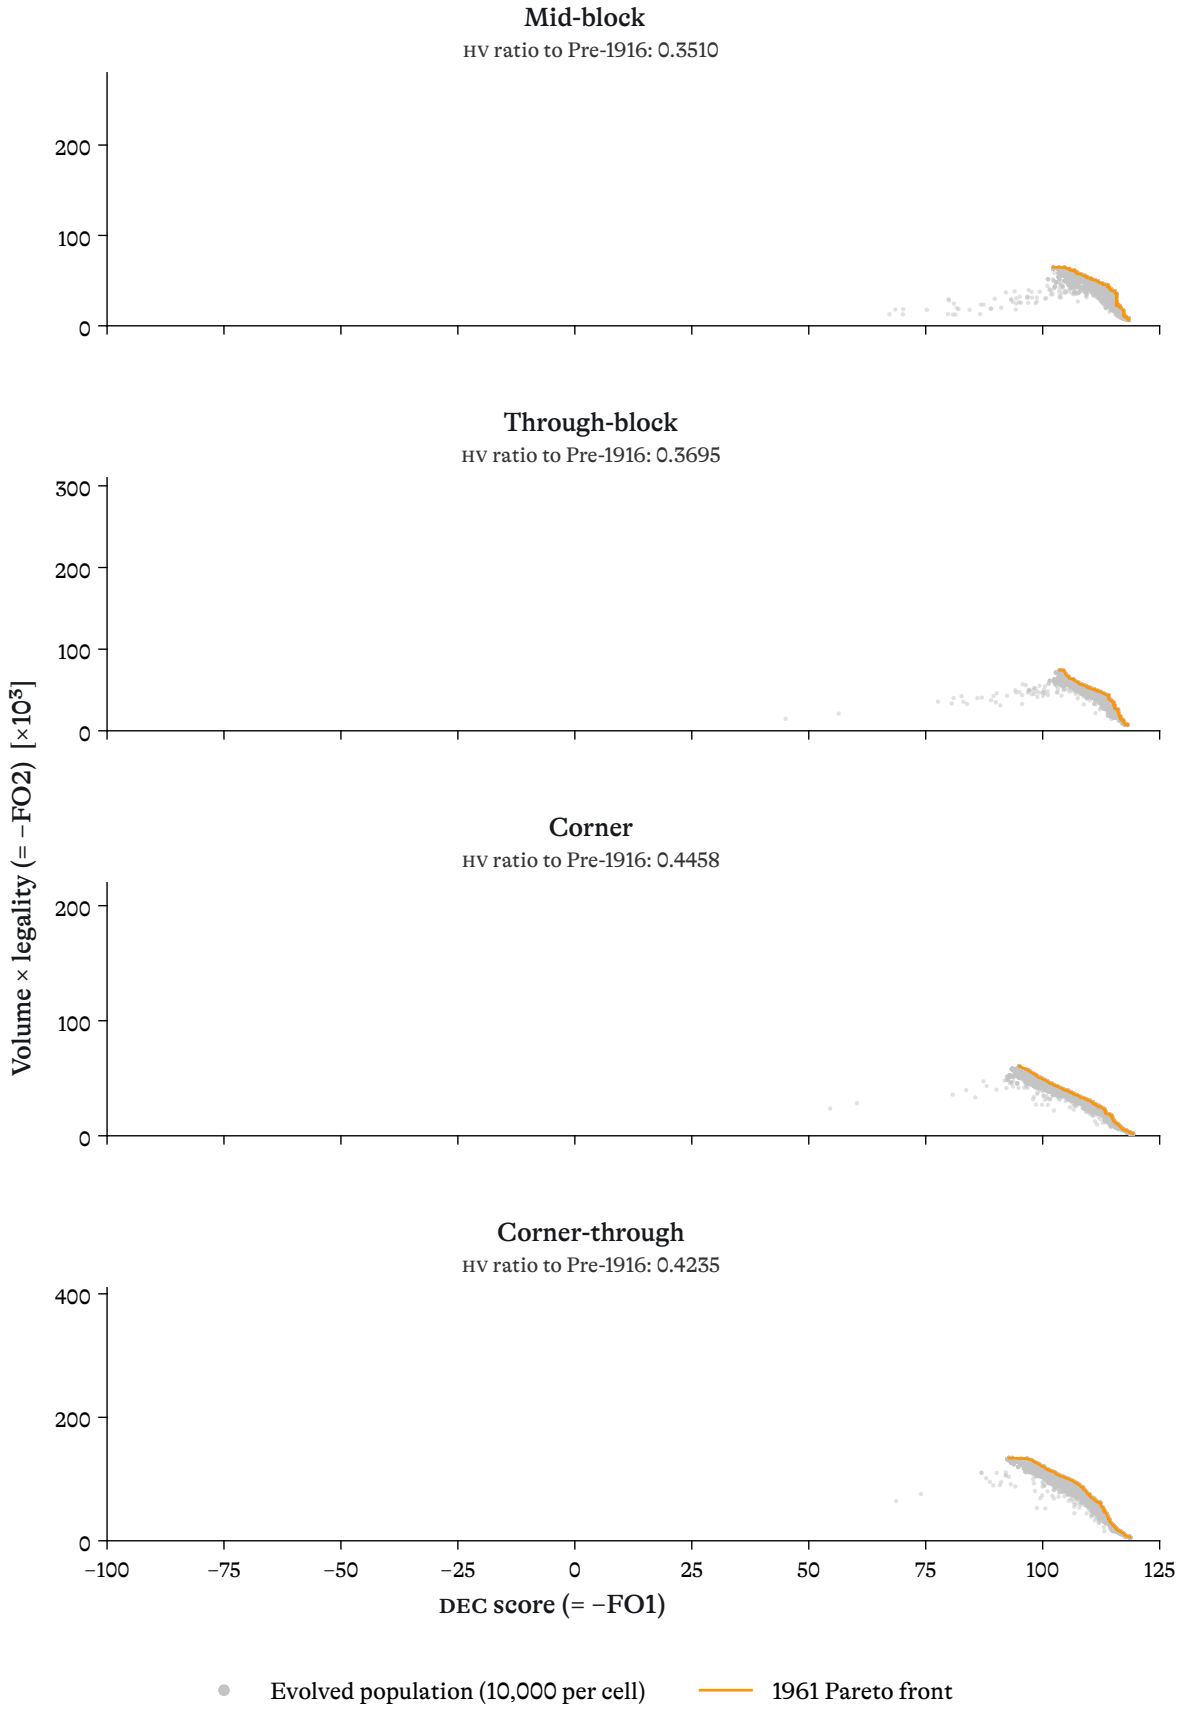


Figure D.5 — 1961: evolved population (light) and Pareto front by typology, with the HV ratio to the pre-1916 baseline noted per panel.

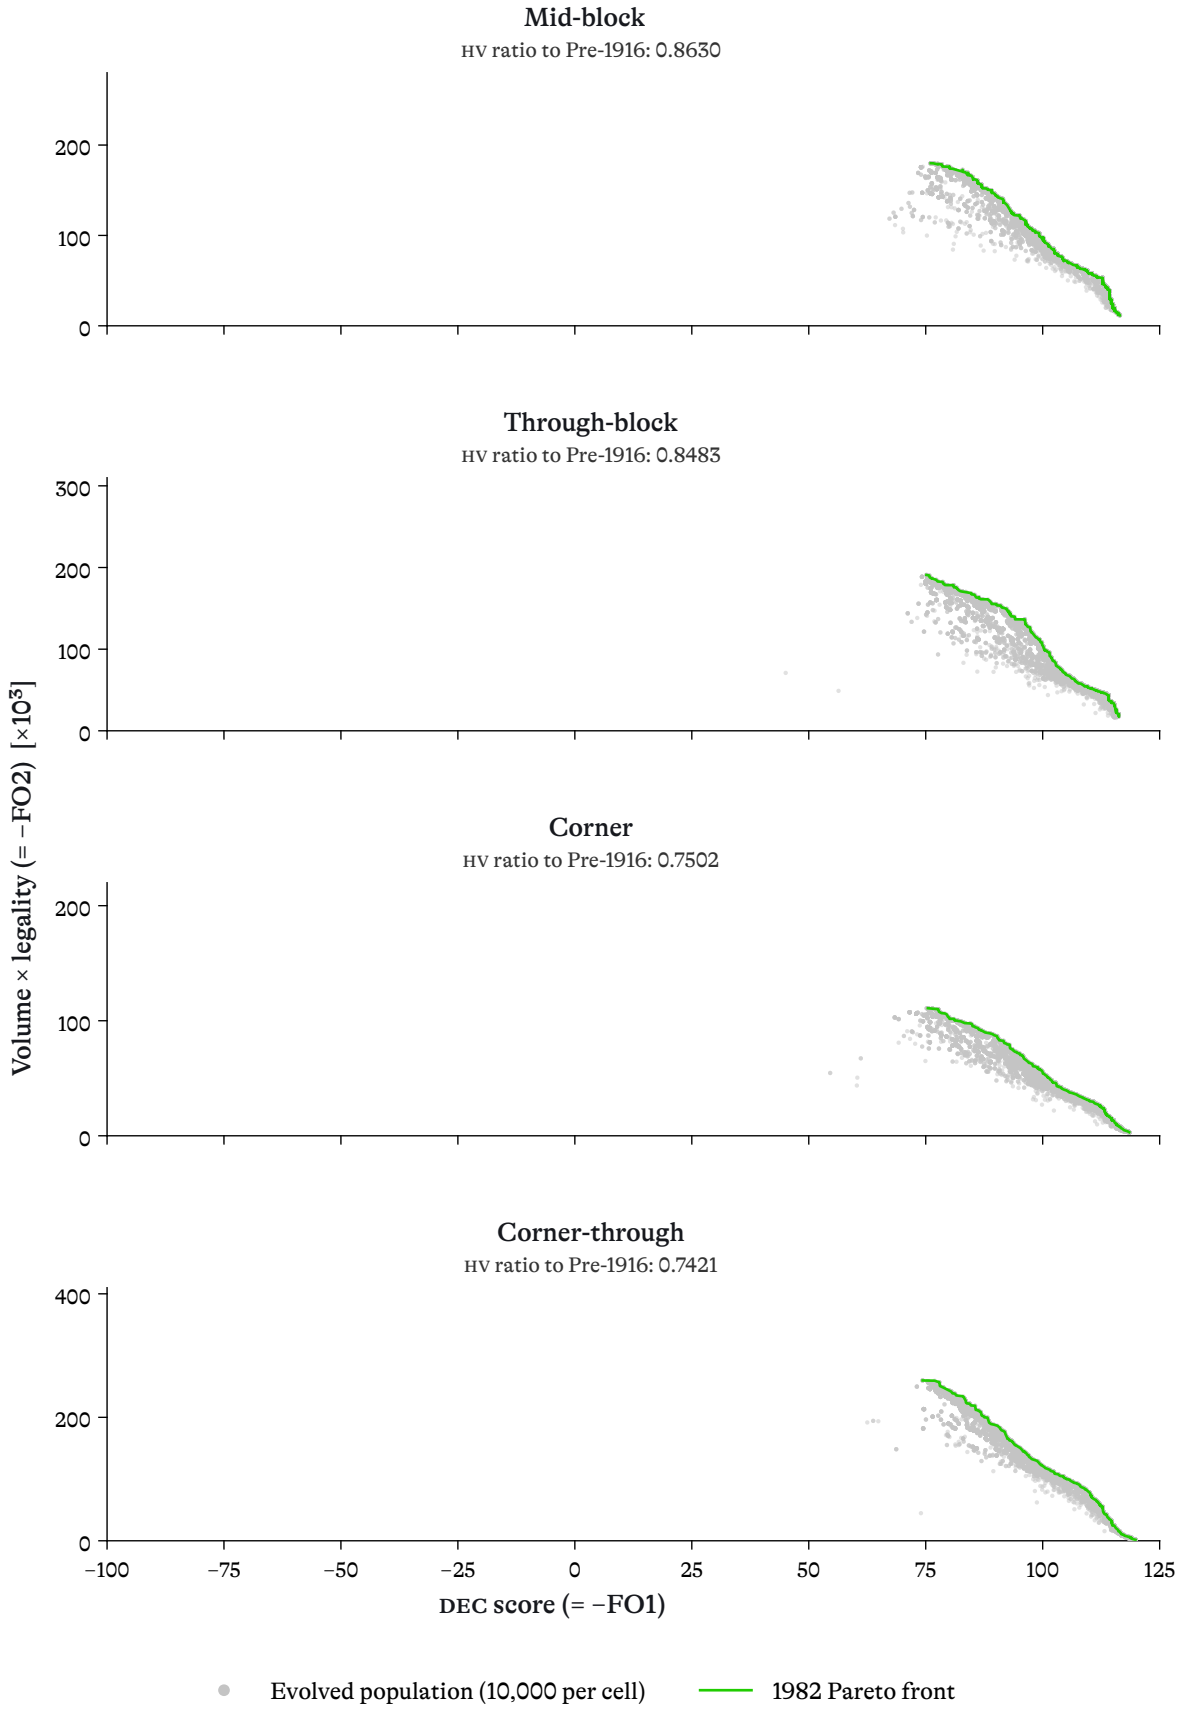


Figure D.6 — 1982: evolved population (light) and Pareto front by typology, with the HV ratio to the pre-1916 baseline noted per panel.

D.4 Full numerical results

Typology\Regime	Pre-1916	1916	1961	1982
Corner	14,738,580.524	8,137,818.951	6,570,714.889	11,056,703.473
Mid-block	21,079,310.099	12,447,610.202	7,399,283.898	18,192,321.010
Through-block	22,658,511.076	10,756,243.897	8,372,331.072	19,220,604.608
Corner-through	34,673,638.452	19,354,809.948	14,683,669.839	25,730,630.213

Table D.1 — Raw HV by regime and typology (objective-space units).

Table D.1 gives the raw HV of every cell in objective-space units (the denominators behind the Table 4.1 ratios). Table D.2 gives the off-front integrity gate behind every H3 statistic: each cell's stratified primary sample was rebuilt genome by genome and re-scored live, and every row had to match its recorded fitness to four decimal places. Fifteen cells pass in full; in the sixteenth, the 1916 through-block, eleven rows were excluded as victims of a Wallacei fill-forward artefact, in which a lagging canvas stamps a stale fitness into the record; their live-measured values leave the cell's coefficient unchanged to three decimals, and all later cells sample phantom-free by construction. Table D.3 collects the per-cell headline metrics.

Typology\Regime	Pre-1916	1916	1961	1982
Corner	242 (0)	185 (0)	128 (0)	208 (0)
Mid-block	258 (0)	227 (0)	136 (0)	249 (0)
Through-block	261 (0)	105 (11)	105 (0)	169 (0)
Corner-through	252 (0)	218 (0)	153 (0)	221 (0)

Table D.2 — Off-front integrity gate: primary sample size per cell, with gate-excluded rows in parentheses.

Regime	Typology	HV ratio	Off-front ρ	On-front ρ	Excluded	Gen → 99%
Pre-1916	Corner	1.0000	+0.977	+0.906	0	92
Pre-1916	Mid-block	1.0000	+0.959	+0.968	0	76
Pre-1916	Through-block	1.0000	+0.941	+0.986	0	91
Pre-1916	Corner-through	1.0000	+0.956	+0.915	0	87
1916	Corner	0.5521	+0.971	+0.928	0	72
1916	Mid-block	0.5905	+0.977	+0.992	0	78
1916	Through-block	0.4747	+0.974	+0.989	11	75
1916	Corner-through	0.5582	+0.982	+0.982	0	79
1961	Corner	0.4458	+0.942	+0.969	0	93
1961	Mid-block	0.3510	+0.962	+0.950	0	88
1961	Through-block	0.3695	+0.970	+0.660	0	92
1961	Corner-through	0.4235	+0.964	+0.995	0	74
1982	Corner	0.7502	+0.868	+0.897	0	78
1982	Mid-block	0.8630	+0.929	+0.826	0	60
1982	Through-block	0.8483	+0.925	+0.994	0	94
1982	Corner-through	0.7421	+0.969	+0.987	0	81

Table D.3 — Per-cell summary: coupling, gate exclusions and convergence, with the HV ratio and the off- and on-front rank coupling given per cell.

Appendix E: Component Structure

This appendix documents the structure of the computational components that constitute the study's instrument, as flowchart diagrams in the standard notation of ISO 5807: rounded terminators, rectangular processes, diamond decisions, parallelograms for data, and double-struck rectangles for predefined processes, calls into another diagrammed component. Nine components are diagrammed here: the phenotype builder; the four regime modules in chronological order, namely the unconstrained pre-1916 baseline, the 1916 height-district envelope, the 1961 prescriptive rule gates, and the 1982 module whose legality is the daylight test itself; the DEC projection; the DEC scorer; the sky-view sampler; and the HV calculator. Charting the full regime family makes the shared module interface visible: identical inputs and audit channel, with only the rule gates differing. Each diagram is drawn from the canonical source used to produce the results of Chapters 3 and 4, and each is followed by a short description. The full, annotated sources of every component are provided in the digital annex accompanying the thesis.

E.1 PhenotypeBuilder: Genome to massing

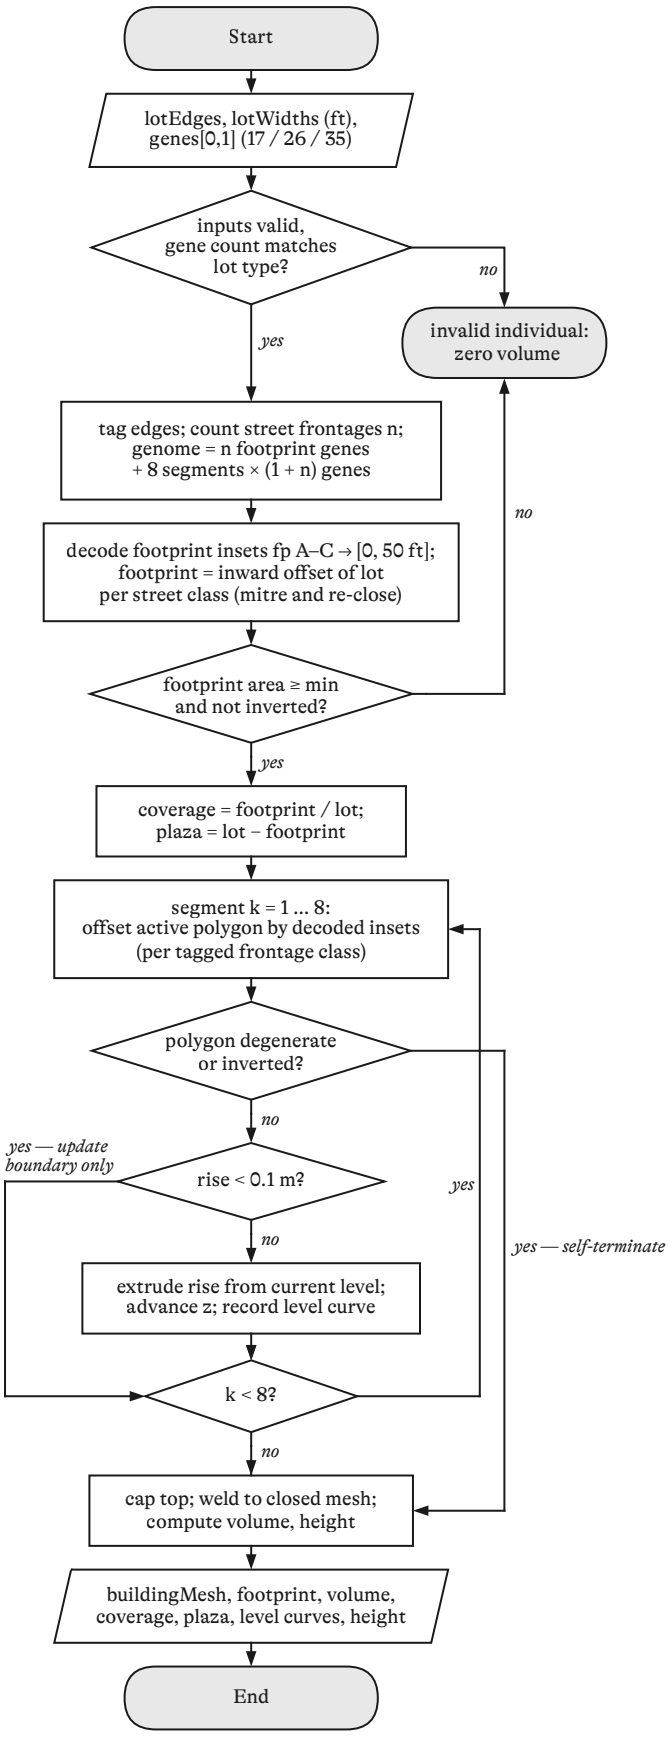


Figure E.1 — PhenotypeBuilder: decoding a fixed-length genome into a closed massing mesh.

The builder decodes a fixed-length vector of real genes, 17, 26 or 35 depending on the number of street frontages, into a massing. The leading genes set a footprint inset per street frontage class, remapped to a 0–50 ft range, so the ground storey may stand back from the lot; the remainder decodes as eight segments of one rise gene plus one inset gene per tagged frontage class. Setbacks are true inward polygon offsets with a mitre-and-reclosure rule, so one decoding serves every lot typology. Degeneracy and polygon-inversion guards self-terminate the stack, and a zero rise updates the active boundary without extruding, which is how the number of visible setbacks becomes an emergent property of the genes. A collapsed footprint returns an invalid, zero-volume individual. The welded, closed output mesh yields volume and height, and coverage ratio and plaza area fall out of the lot–footprint distinction as free descriptive statistics. With all footprint genes at zero the builder reproduces the lot-boundary decoding exactly, its regression gate.

E.2 RegimePre1916: The unconstrained baseline

The unconstrained baseline charted so the family interface is explicit. The module validates its inputs and applies no rule: pre-1916 Manhattan had no comprehensive height, setback, floor-area or coverage control, so every valid massing is legal at score 1.0 and the two objectives reduce to raw daylight and volume. Same inputs, same audit channel, no gates.

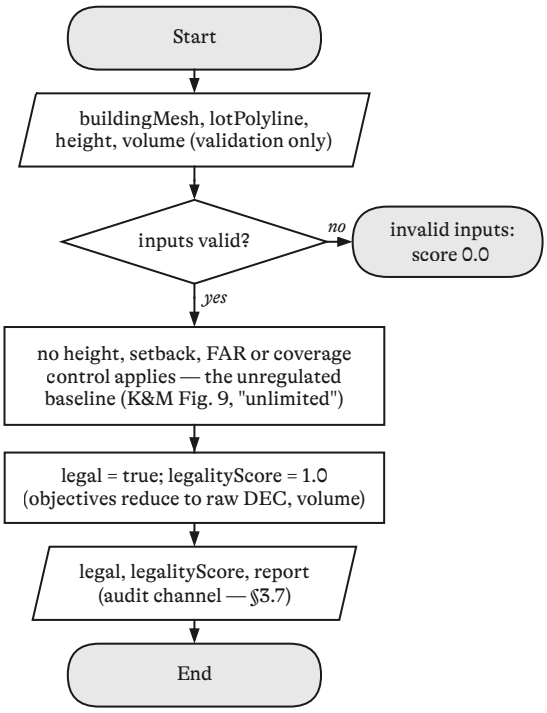


Figure E.2 — RegimePre1916: input validation and the null rule set. The unregulated baseline against which the three regulated regimes are measured.

E.3 Regime1916: Height districts and the setback envelope

The 1916 module reconstructs the height-district envelope as calibrated working values (Appendix A.2): a street wall of twice the mapped street width, and a sky plane rising four-to-one behind the street line. Each storey's level curve is sampled against the envelope of every street frontage; a storey that exceeds it may still pass under the tower exemption: footprint at most a quarter of the lot, everywhere at least seventy-five feet from each street centreline. Legality requires every storey to pass, and the graded score is the passing fraction, reported as the audit channel of §3.7, since the decoder's clipping means a non-compliant storey never reaches the population.

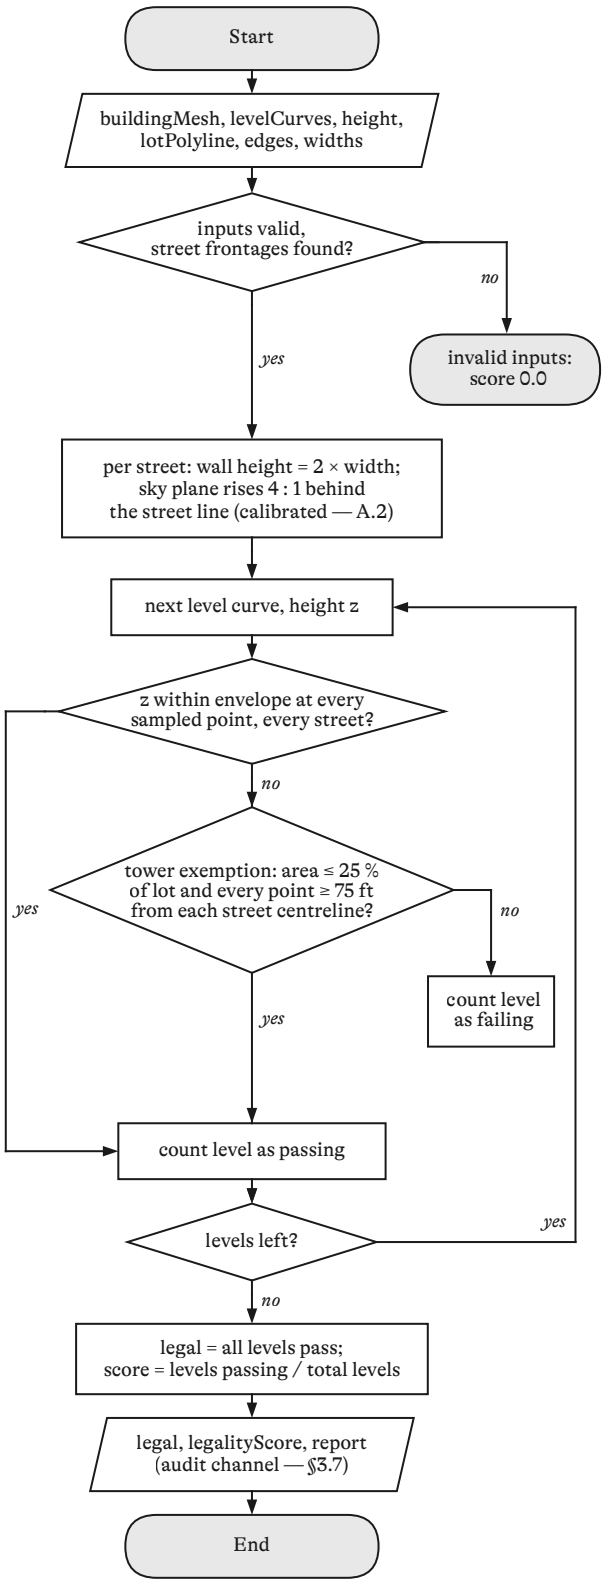


Figure E.3 — Regime1916: per-storey envelope test with the tower exemption, and the passing-fraction legality score.

E.4 Regime1961: The 1961 prescriptive rule gates

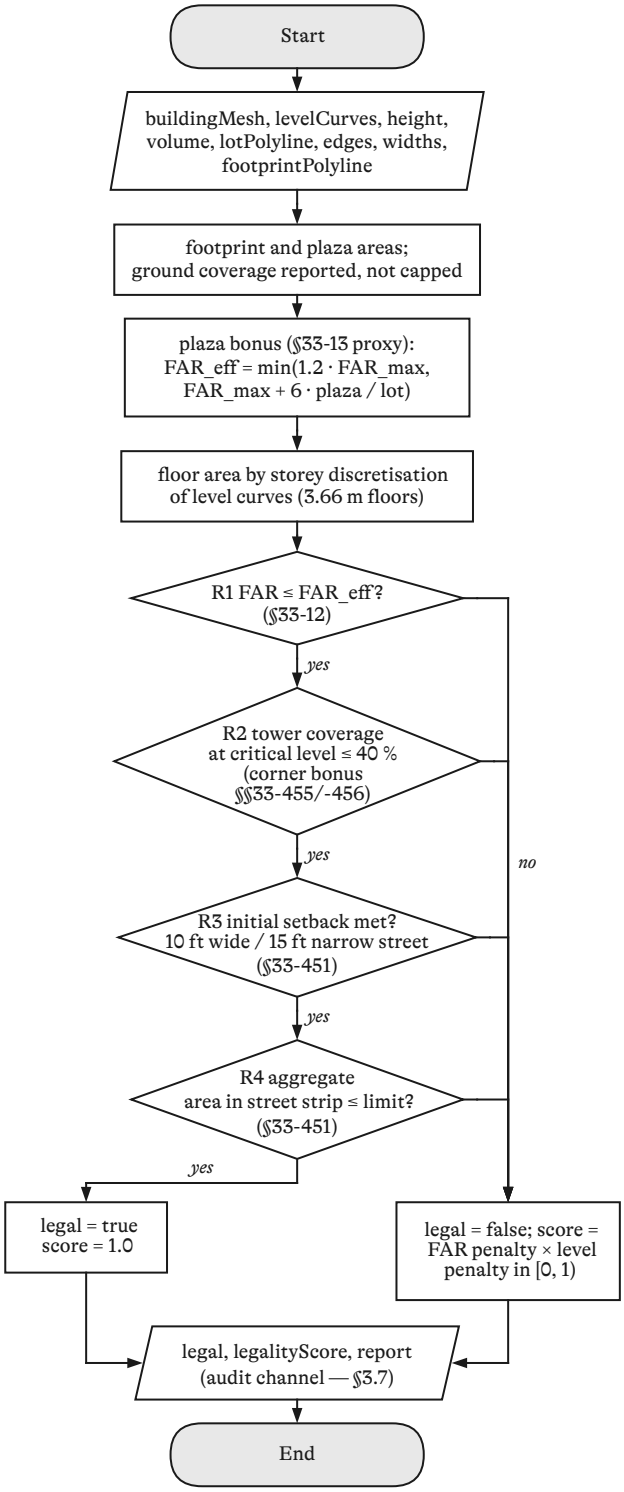


Figure E.4 — Regime1961: floor-area ratio, tower coverage, setback and street-strip rules with the plaza bonus.

The 1961 module is diagrammed as the representative of the prescriptive constraint components, being the most rule-dense. It consumes the common regime interface, namely massing, level curves, lot geometry, tagged edges and widths, and computes floor area by storey discretisation of the level curves at the representative 3.66 m floor. The plaza bonus is a proxy for §33-13: the effective ceiling is the base ratio of fifteen raised by six times the plaza fraction, capped at twenty per cent over base, so a recessed footprint earns floor area rather than obeying a coverage rule; ground coverage is reported, never constrained. Four rules are tested in sequence: the floor-area ceiling of §33-12, tower coverage of forty per cent at the critical level with the corner alternates of §§33-455 and 33-456, the initial setback of ten or fifteen feet by street width, and the aggregate street-strip area of §33-451. A legal massing scores one; an illegal one a graded product of the ratio shortfalls, the audit channel of §3.7, since the decoder clips phenotypes to the envelope.

E.5 Regime1982: Legality as daylight performance

The 1982 module is deliberately thin, because under this regime daylight performance is the law. The DEC score and per-street scores arrive from DECScorer (Figure E.7); the module owns only the statutory thresholds of §81-274(i), seventy-five overall and sixty-six on every street, and the soft penalty 1 − max(shortfall ratios) that grades near-misses. It recomputes the thresholds rather than trusting the scorer's flag, and cross-checks the two so the regime layer stays self-documenting.

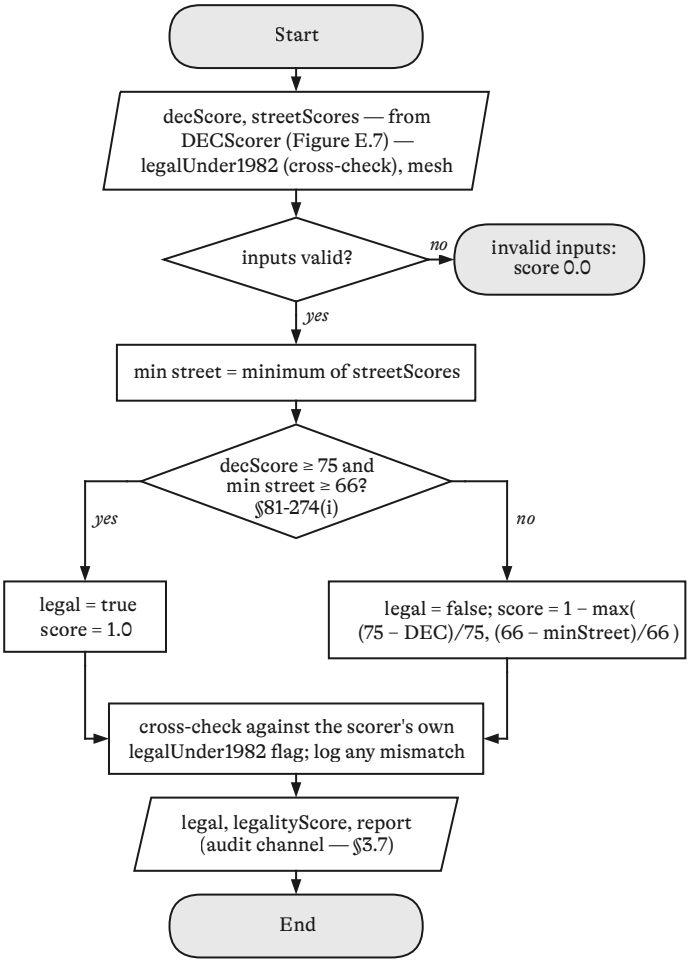


Figure E.5 — Regime1982: the §81-274(i) thresholds over the scorer's outputs, with the graded shortfall penalty.

E.6 DECProjection: The rectilinear chart projection

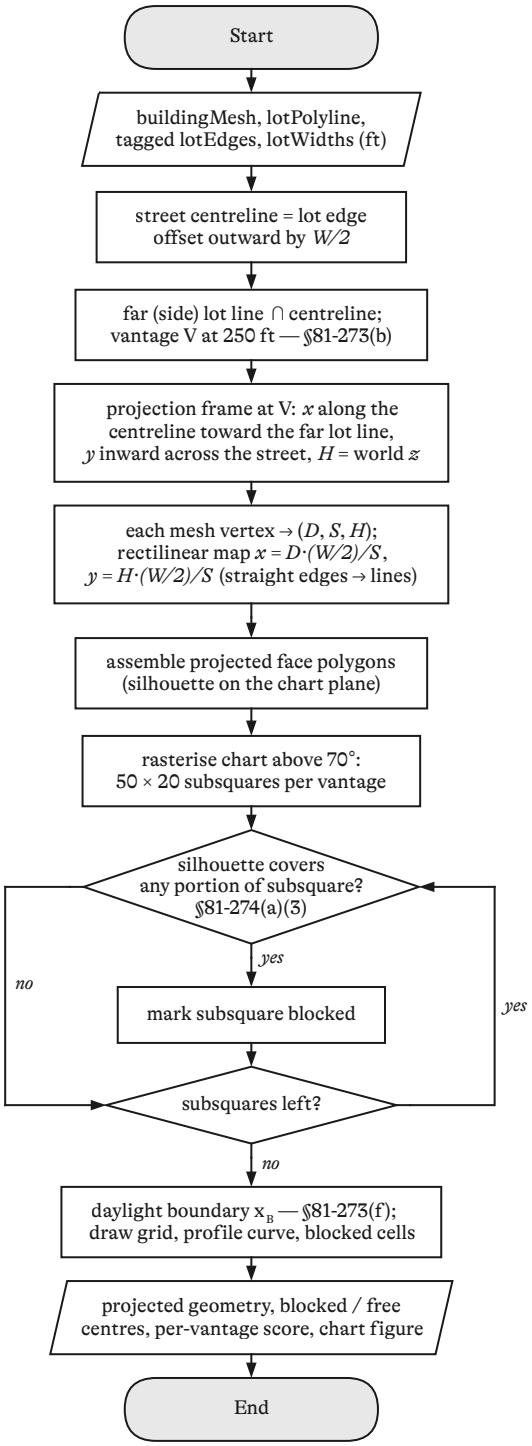


Figure E.6 — DECProjection: vantage construction, rectilinear projection and any-portion rasterisation for one vantage point.

The projection is the single-vantage twin of the scorer, used for validation figures and debugging. The vantage point stands on the street centreline, 250 feet from the centreline's intersection with the far side lot line, per the plotting rules of §81-273(b). A projection frame is erected there: x along the centreline toward the far lot line, y inward across the street, height from world z. Every mesh vertex maps to the chart plane by the rectilinear transformation of Fishhaut, Weingarten and Neiger (1984, pp. 185–186):

$$x = \frac{D \cdot \left(\frac{W}{2}\right)}{S} \tag{E.1}$$

$$y = \frac{H \cdot \left(\frac{W}{2}\right)}{S} \tag{E.2}$$

where, for the vertex under projection, D is its distance from the vantage point along the street axis (ft); S is its sagittal distance from the vantage, the offset component perpendicular to the street axis (ft, $S > 0$; vertices at or behind the vantage plane are discarded); H is its height above the street plane (ft); and W is the street width fed to the scorer (ft), so that $W/2$ is the vantage's centreline offset and the chart's scale factor. The transformation sends straight building edges to straight lines.

The chart above seventy degrees is rasterised into fifty by twenty subsquares, and a subsquare is charged when the projected silhouette covers any portion of it, the canonical rule of §81-274(a)(3), implemented by a predicate shared verbatim with the scorer so the two instruments are in parity. The component draws the grid, the daylight boundary of §81-273(f), the profile curve and the blocked cells, the visual form of the validation of §3.9.

E.7 DECScorer: The production scorer

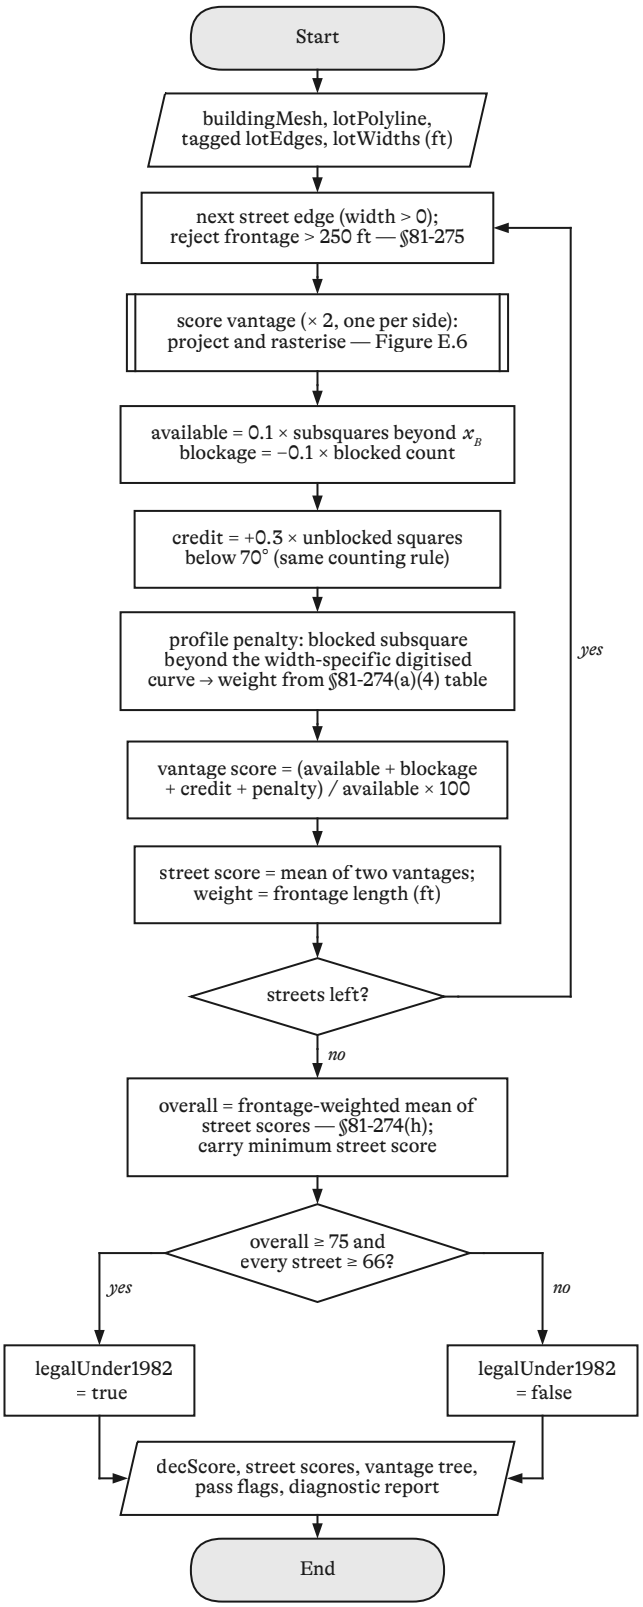


Figure E.7 — DECScorer: per-vantage scoring arithmetic and the frontage-weighted aggregation of §81-274(h).

The scorer is the production instrument, identical in both evaluations, a pure function from massing, lot and street widths to a score. Each street frontage receives two vantage points, one per side, each scored through the projection pipeline of Figure E.6. Per vantage, available daylight is one tenth of a square per subsquare beyond the daylight boundary; blockage subtracts one tenth per charged subsquare; unblocked squares below seventy degrees earn a credit of three tenths under the same counting rule; and blocked subsquares lying beyond the width-specific profile curve, a nineteen-point table digitised from the official chart graphics, draw the penalty weights of §81-274(a)(4). The vantage score is the remaining fraction of available daylight, expressed as a percentage. Street scores are the mean of their two vantages; the overall score is their frontage-length-weighted mean per §81-274(h), with the minimum street score carried separately against the sixty-six per cent test. Frontages over 250 feet are rejected as out of scope (§81-275), and both thresholds together set the legality flag consumed by the 1982 module.

E.8 CenterlineSVF: The independent openness sampler

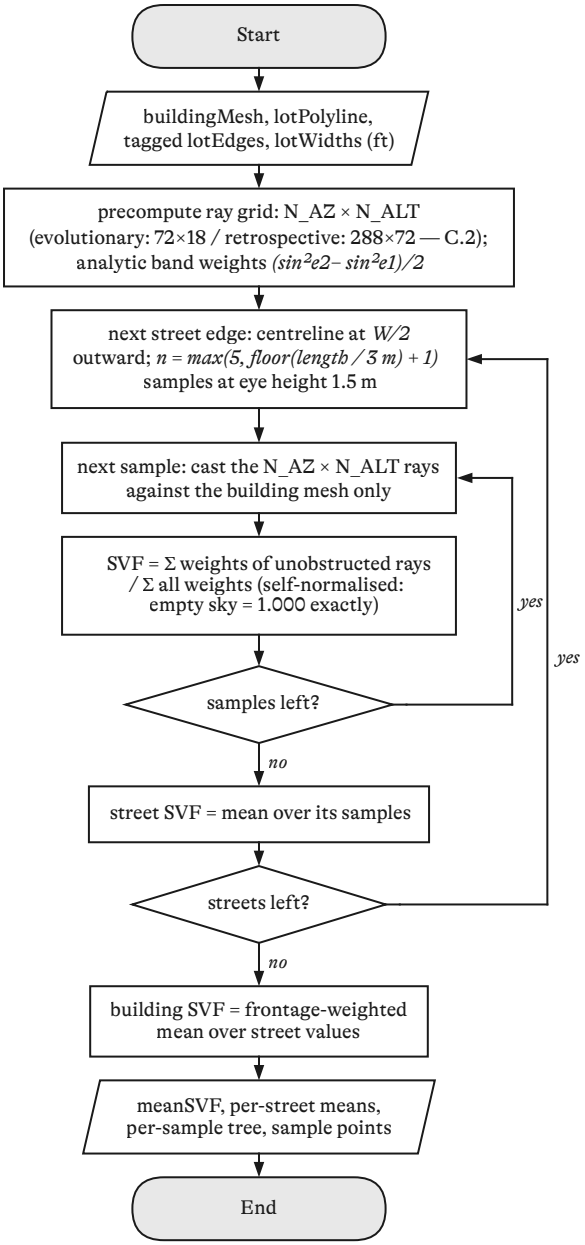


Figure E.8 — CenterlineSVF: hemispheric ray sampling along the street centreline at pedestrian eye height.

The sampler supplies the thesis's independent measure of street-level openness; it is never an objective and is computed only after selection has finished. The upward hemisphere is discretised on an azimuth-by-elevation ray grid: 72 by 18 (1,296 rays) for the evolutionary evaluation and 288 by 72 (20,736) for the retrospective evaluation, the resolution policy and its convergence evidence set out at C.2. The band weights are the exact analytic $(\sin^2 e_2 - \sin^2 e_1)/2$, and the sky view factor at a sample is the weight fraction of rays that clear the building; the measure is self-normalised, so an empty sky

returns exactly one at any grid resolution. Samples stand on each street's centreline at a pedestrian eye height of 1.5 metres, spaced at three metres with at least five per frontage so short frontages are not under-sampled. Rays test the building mesh alone, the building-only convention of §3.5 that mirrors what the Chart itself is allowed to see. Street values are the mean over their samples; the building value is the frontage-weighted mean over its streets; and the per-sample tree supplies the along-frontage profiles the fixed vantage points of the Chart structurally cannot record.

E.9 HypervolumeCalc: Measuring the design spaces

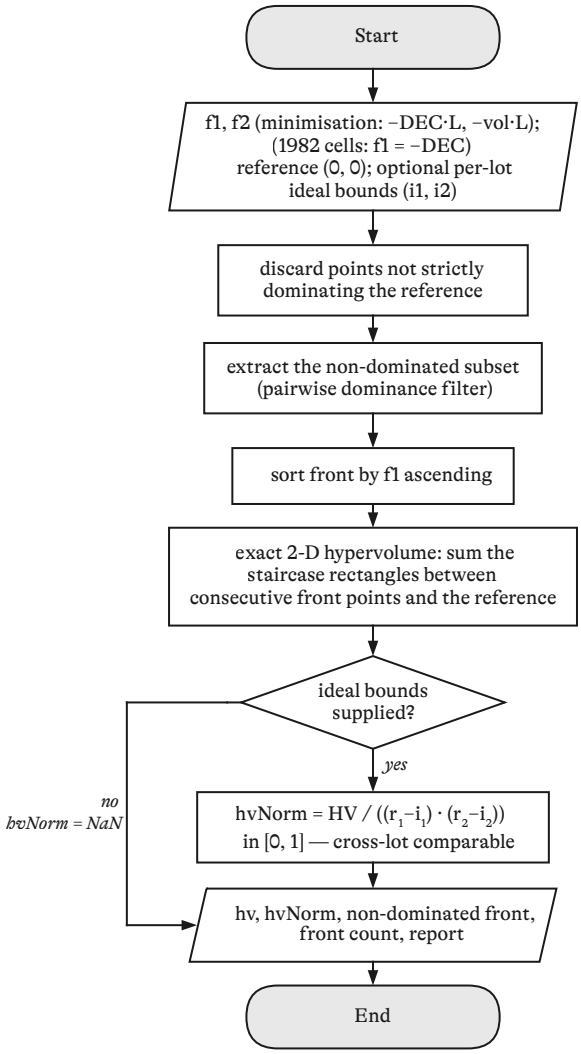


Figure E.9 — HypervolumeCalc: exact two-dimensional HV of a Pareto front against the reference point.

The calculator measures the size of each evolved design space as the area of objective space its Pareto front dominates, following Guerreiro, Fonseca and Paquete (2021). Both objectives arrive in minimisation form, namely the negated products of daylight score and legality, and of volume and legality. The reference point is the origin, common to all sixteen cells so the measurements are commensurable. The component accepts a whole population of any generational mix: points that do not strictly dominate the reference are discarded, the non-dominated subset is extracted by a pairwise dominance filter, and the front is sorted by the first objective. In two dimensions the HV is then exact and cheap, the sum of the staircase rectangles between consecutive front points and the reference, with no approximation involved. The ratios of §4.2 are formed downstream, each regime's raw HV divided by the pre-1916 value on the same lot (Table D.1). The analytic self-test of Appendix C.3, a three-point front of known HV, gates the component before use.

The remaining components follow the same single-responsibility pattern and are not diagrammed. LotLibrary defines the four canonical lot typologies (Figure 3.1) with tagged frontage edges and widths, held fixed across regimes; the lot pre-processor emits the canonical geometric frame; and the instrument test rigs, DECTestRig and SVFTestRig, generate canonical geometry and closed-form expected values for the verification programme of §3.9. The parameters of every regime module are tabulated in Appendix A, the scoring algorithms in Appendix C, and the full annotated source of all components in the digital annex.