Walkthrough of the Canonical Figures

Gil Raitses

April 24, 2026

Abstract

This document walks through the canonical figures the routing substrate has produced over the past four months. It is structured as a self-contained surface for Jen Schwarz to engage with at her pace. Each figure occupies a full page with a structured caption (title, overview, per-panel descriptions) followed by a structured discussion (remarks, claims, boundaries, limitations, things to improve, questions for Jen) on the subsequent pages.

The substrate models pedestrian routing on the Midtown Manhattan corridor (40th to 59th Streets, Lexington to 8th Avenues) as a graph problem with two operative measurements: a camera-derived stress field (demand side, per-block, locked in a project decision record) and a seven-layer cross-street capacity model (supply side, per-edge, derived from twenty-five public NYC Open Data sources). The routing layer combines them as cost equals stress over capacity. The substrate's core empirical findings include (i) the routing layer produces strictly dominant alternative paths when given access to the second-order capacity field; (ii) the substrate exhibits an allosteric soft-mode redistribution under static perturbation, with zero leakage across the corridor boundary; (iii) the supply and demand layers are orthogonal at the per-edge level (Pearson r = -0.014); (iv) the lived-knowledge audit aligns at 8 of 12 reference blocks, with zero misalignments.

The synthesis question at the end of this document asks whether the substrate's behavior is consistent with operating near a rigidity transition in your formalism, and what control parameter and response amplitude metric you would suggest if so. Two recent papers identified in the past week (Sridhar et al. on Allocentric Flocking and Sane et al. on Generalization at the Edge of Stability) are referenced in the per-figure discussions as cross-domain anchors that make the same methodological move in different fields.

Preface

This document walks through the canonical figures the routing substrate has produced over the past four months, with the structured remarks, claims, boundaries, limitations, things to improve, and questions for you that I have accumulated about each one. The goal is to give you a single self-contained surface that captures the substrate's current state and the open conceptual questions I would value your input on.

The figures are grouped by purpose. Section 1 covers the routing demonstration figures (the comparison between shortest-path and stress-aware routing, the generalization of the comparison across destinations, and the multi-scale hierarchical variant). Section 2 covers the substrate's underlying second-order constraint field (the camera-derived stress heatmap, the cross-street capacity overlay, the per-edge route alternatives). Section 3 covers the spectral analysis of the corridor graph operator (the Laplacian eigenvectors, the centrality measures, the stress-flow coupling, the boundary flux). Section 4 covers the validation figures (the lived-knowledge audit, the orthogonality check between supply and demand). Section 5 covers extended scenarios (time-varying capacity, spillover modes).

The synthesis section at the end consolidates the open conceptual questions across all figures into one

synthesis question for you. The references section at the very end lists all the related literature cited in the per-figure discussions, organized by domain.

1. Routing Demonstration

1.1 Shortest vs. stress-aware routing

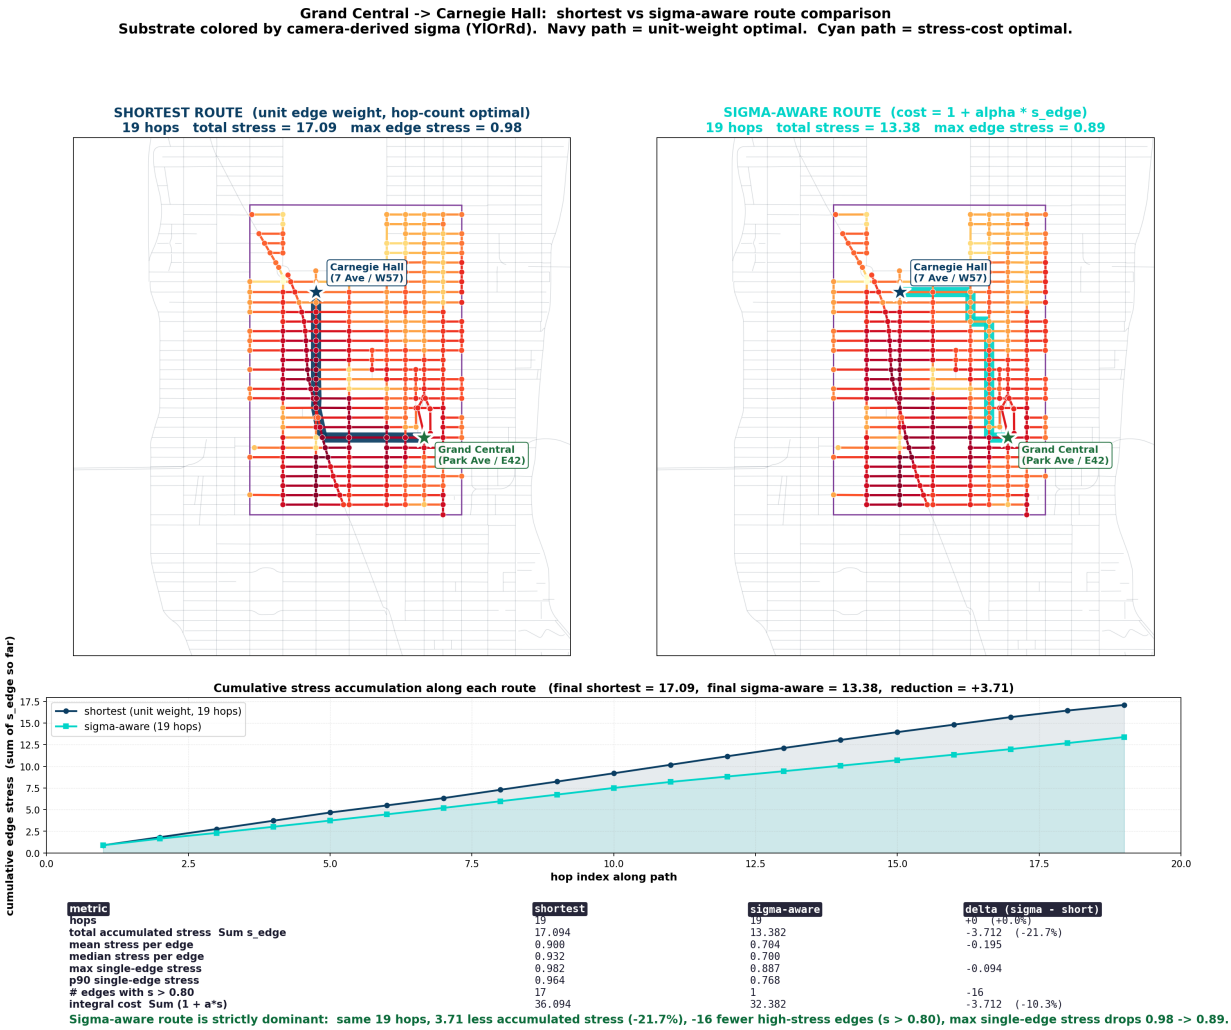

Figure 1. Shortest vs. stress-aware routing for Grand Central to Carnegie Hall.
The figure compares two routes between the same origin and destination on the Midtown corridor graph. Both routes are nineteen hops. The substrate is colored by camera-derived stress (YlOrRd, pale to deep red as stress rises). Two route maps sit side by side; a cumulative-stress chart sits below; a metrics table sits at the foot.
Top left, shortest route. The unit-edge-weight optimal path under standard A* routing. Total accumulated stress is 17.09 and the maximum single-edge stress is 0.98. The route runs along high-stress avenues without measuring them.
Top right, sigma-aware route. The same A* pathfinding under cost $1 + \alpha s_{\text{edge}}$. Total accumulated stress is 13.38 and the maximum single-edge stress is 0.89. Same hop count; sixteen fewer high-stress edges; max single-edge stress drops by 0.09.
Bottom, cumulative stress along each route. Both curves start at the origin and accumulate per-edge stress with each hop. The sigma-aware route stays consistently below the shortest route, with a divergence that grows monotonically along the path.
Foot, metrics table. Hops, total accumulated stress, mean stress per edge, max single-edge stress, count of edges above 0.80, integral cost. The sigma-aware route is strictly dominant on every measured criterion at the same path length.

Remarks. This is the demonstration's anchor. It shows the two routes the routing layer can produce given access to either the first-order graph alone or the first-order graph plus the second-order capacity field. The contrast is unambiguous: same task, same hop count, very different second-order load.

Claims. The routing layer can produce a strictly dominant alternative when given access to the second-order constraint field. The dominance is in measured stress accumulation, not in inferred user comfort. The figure does not claim a behavioral preference; it claims a measurable substrate-level difference.

Boundaries. The figure is for one origin-destination pair (Grand Central to Carnegie Hall). The substrate covers the Midtown corridor (40th to 59th Streets, Lexington to 8th Avenues). The shortest path is the unit-weight A* optimal; alternative shortest-path tie-breaks may produce slightly different specific routes but the same overall stress profile.

Limitations and caveats. The stress field is camera-derived, not behaviorally validated. The 22% reduction in accumulated stress is a substrate-level number; whether a real pedestrian would experience the second route as 22% less stressful is not measured here. The cumulative-stress curve aggregates per-edge stress without weighting for path length, time of day, or pedestrian profile.

Things to improve. Add a third panel: the stress-aware route under a time-varying capacity field (peak-hour subway egress, active scaffolding sheds), to show the recommendation shifting with the field. The current figure is static.

Questions for Jen. The cumulative-stress curve has a structural feature I find interesting: the divergence between the two routes is not concentrated at one bottleneck edge but accumulates monotonically along the path. Does this match the soft-mode picture you would expect from a system operating near a rigidity transition, where a small constraint perturbation produces redistributed (rather than localized) response? I read the monotonic divergence as evidence of the substrate operating in a regime where local edge cost differences integrate over path length to produce path-level differences, but I am not sure if that is the right physical intuition.

1.2 Routes generalize across destinations

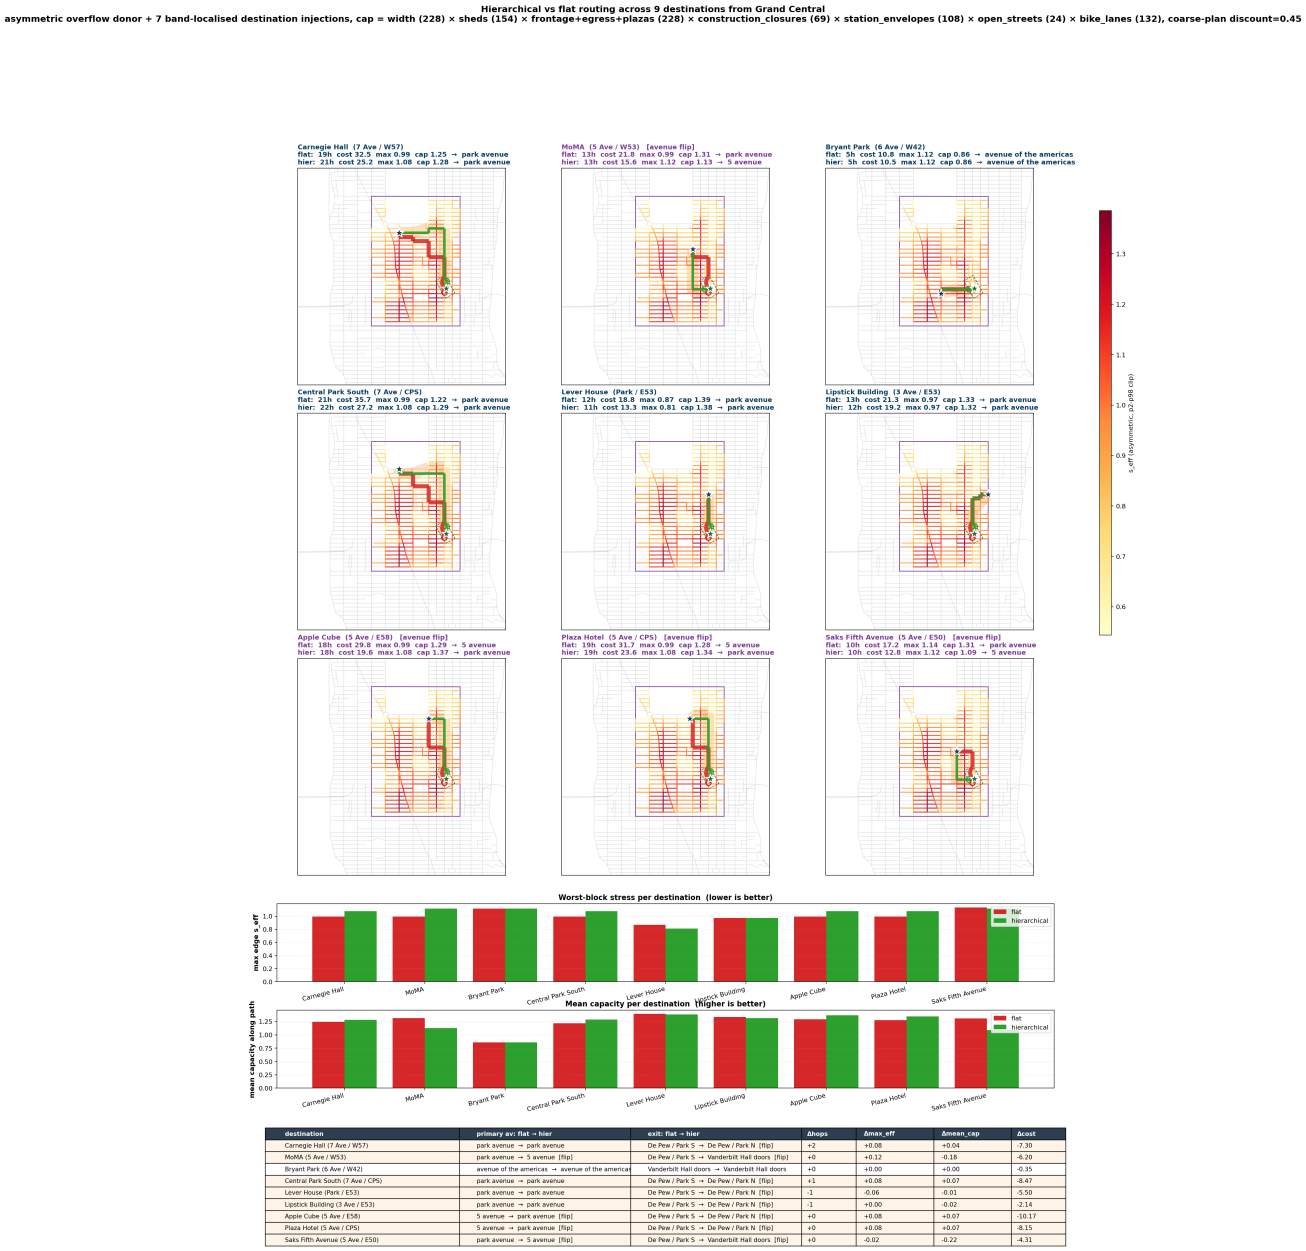

destination	primary av: flat → hier	exit: flat → hier	Δhops	Δmax_eff	Δmean_cap	Δcost
Carnegie Hall (7 Ave / W57)	park avenue → park avenue	De Pew / Park S → De Pew / Park N [flip]	+2	+0.08	+0.04	-7.30
MoMA (5 Ave / W53)	park avenue → 5 avenue [flip]	De Pew / Park S → Vanderbilt Hall doors [flip]	+0	+0.12	-0.18	-6.20
Bryant Park (6 Ave / W42)	avenue of the americas → avenue of the americas	Vanderbilt Hall doors → Vanderbilt Hall doors	+0	+0.00	+0.00	-0.35
Central Park South (7 Ave / CPS)	park avenue → park avenue	De Pew / Park S → De Pew / Park N [flip]	+1	+0.08	+0.07	-8.47
Lever House (Park / E53)	park avenue → park avenue	De Pew / Park S → De Pew / Park N [flip]	-1	-0.06	-0.01	-5.50
Lipstick Building (3 Ave / E53)	park avenue → park avenue	De Pew / Park S → De Pew / Park N [flip]	-1	+0.00	-0.02	-2.14
Apple Cube (5 Ave / E58)	5 avenue → park avenue [flip]	De Pew / Park S → De Pew / Park N [flip]	+0	+0.08	+0.07	-10.17
Plaza Hotel (5 Ave / CPS)	5 avenue → park avenue [flip]	De Pew / Park S → De Pew / Park N [flip]	+0	+0.08	+0.07	-8.15
Saks Fifth Avenue (5 Ave / E50)	park avenue → 5 avenue [flip]	De Pew / Park S → Vanderbilt Hall doors [flip]	+0	-0.02	-0.22	-4.31

Figure 2. Stress-aware routing across nine destinations from Grand Central.
The figure shows the substrate's recommended routes for nine different destinations, with both flat (red) and hierarchical (green) routing modes overlaid on each panel. The full seven-layer capacity model is active (sidewalk width, sheds, frontage load, construction closures, station envelopes, open streets, bike lanes). Each panel is one destination, with the metrics table at the foot summarizing the comparison.
Nine destination panels. Carnegie Hall, MoMA, Bryant Park, Central Park South, Lever House, Lipstick Building, Apple Cube, Plaza Hotel, Saks Fifth Avenue. Each panel shows the corridor substrate colored by combined effective cost (s_{eff}, p2–p98 clipped, navy to deep red). Red overlay is the flat-routing recommendation; green overlay is the hierarchical-routing recommendation. Annotations on each panel give hops, total cost, max edge cost, and mean capacity for each mode.
Foot, metrics summary table. Nine destinations × four columns: primary avenue (flat → hierarchical), exit (flat → hierarchical), Δhops, Δmax_eff, Δmean_cap, Δcost. The hierarchical mode produces substantially lower total cost on most destinations (negative Δcost values) at the price of a small max-edge-cost increase on some.

Remarks. The contrast generalizes across destinations. The flat route differs from the hierarchical route on eight of the nine destinations. The metric trade-off is consistent: hierarchical routing chooses an avenue spine that lowers total cost while accepting a slightly higher max-edge cost on the chosen avenue.

Claims. The substrate's recommendations are not artifacts of one origin-destination pair. Same routing layer, same capacity model, nine different destinations, and the structural argument holds. The flat-vs-hierarchical contrast also holds: hierarchical routing is doing real work, not just stylistic re-routing.

Boundaries. Nine destinations is a sampling, not a census. All destinations are within the Midtown corridor; the substrate has not been tested on destinations outside the 40th–59th Street range. The capacity model is the seven-layer combined version; the contrast may behave differently if individual layers are removed.

Limitations and caveats. The hierarchical mode applies a coarse-tile-plan discount (factor 0.45) to edges that lie on the planned avenue. This is a heuristic, not a derived parameter; alternative discount values produce different specific routes (though the qualitative picture holds across a range tested). The comparison favors hierarchical routing partly because of this discount.

Things to improve. Sample destinations outside the corridor (e.g., Penn Station to Carnegie Hall) to test whether the substrate generalizes beyond its training-equivalent geographic envelope. Add an ablation: same nine destinations with each capacity layer removed in turn, to identify which layers are doing the work for the contrast.

Questions for Jen. The hierarchical mode reads to me as a coarse-grained scaffolding operation: it imposes a low-frequency structure (which avenue to use) on top of which the high-frequency edge-by-edge routing happens. This feels analogous to the multi-scale moves your group has made in different contexts. Does the coarse-grained / fine-grained decomposition here look like the same kind of scale-separation move you would deploy in a physical system, or is it doing something else operationally?

1.3 Hierarchical multi-scale routing

metric	FLAT	HIERARCHICAL	delta (hier - flat)
hops	18	19	+1
total s_eff	13.716	14.012	+0.296
max edge s_eff	1.043	0.925	-0.118
mean capacity	1.122	1.350	+0.228
total cost (effective)	30.453	23.870	-6.583
total cost (undiscounted)	30.453	31.677	+1.224
first exit chosen	east 43 street - david ben-gurion place|vanderbilt avenue	de pew place|east 45 street|park avenue	—
first exit label	Vanderbilt Hall doors	De Pew / Park N	—
coarse plan	vanderbilt avenue band=9 -> park avenue band=10 -> park avenu	(applied as edge discount)	19 tiles, cost 37.84

Figure 3. Flat vs. hierarchical routing on the asymmetric-spillover field.
The figure isolates the flat-vs-hierarchical comparison for a single destination (Grand Central to Carnegie Hall) under the asymmetric-spillover capacity model with a coarse-tile-plan discount of 0.45. Two route maps sit side by side; an avenue-dwell bar chart sits below; a cumulative-stress chart sits next; a metrics table sits at the foot.
Top left, flat routing. 18 hops, total $s_{\text{eff}} = 13.72$, max edge = 1.04, total cost = 30.45. The route runs primarily along 5th Avenue and exits at East 43rd / Vanderbilt.
Top right, hierarchical routing. 19 hops, total $s_{\text{eff}} = 14.01$, max edge = 0.925, base cost = 31.68, discounted cost = 23.87. The route runs primarily along Park Avenue and exits at De Pew Place / East 45th / Park N.
Middle, avenue-dwell bar chart. Edges spent per avenue or cross-street. The flat mode commits 12 edges to 5th Avenue; the hierarchical mode commits 12 edges to Park Avenue, 6 to a cross-street between, and 1 to an interior connector.
Below, cumulative effective stress along each path. Two curves; the hierarchical curve sits slightly above the flat curve in raw s_{eff} but well below in discounted cost (the green-shaded region in the right portion of the chart shows the discount applied).
Foot, metrics table. Hops, total s_{eff}, max edge s_{eff}, mean capacity, total cost (effective), total cost (undiscounted), first exit chosen, first exit label, coarse plan applied. Hierarchical routing produces a lower discounted total cost at the price of one extra hop and a slightly higher raw stress accumulation.

Remarks. This figure isolates the multi-scale move. Flat routing optimizes hop-by-hop; hierarchical routing biases path choice toward avenues with second-order capacity headroom (Park Avenue here, vs. 5th Avenue under flat routing) by applying a coarse-tile-plan discount.

Claims. The hierarchical layer is doing real work. The flat-vs-hierarchical comparison shows two structurally different paths (different avenue spine, different exit point, different per-block stress profile), not two cosmetic variations. The avenue-dwell chart makes the structural difference visible: 5th vs. Park is the operative choice.

Boundaries. The coarse-tile-plan discount of 0.45 is the canonical setting; alternative discounts in the range 0.3–0.6 produce qualitatively similar pictures. The hierarchical layer operates on avenue-band membership, not arbitrary path features; the avenue-band decomposition is corridor-specific.

Limitations and caveats. The hierarchical mode adds one hop to this specific route. The trade-off is hop count vs. accumulated stress; a pedestrian who values minimum hop count over minimum stress would not benefit from this mode. The discount is applied as a multiplicative factor on edge cost, not as a derived parameter; it is a heuristic.

Things to improve. Replace the heuristic discount with a derived parameter from the substrate's local-vs-global capacity ratio (the ratio of mean per-edge capacity along the chosen avenue to the corridor-mean capacity). This would make the hierarchical layer's behavior continuous with the rest of the model rather than an inserted scaffold.

Questions for Jen. The hierarchical layer feels like a renormalization-group move: coarse-grain to one decision (which avenue), then run the fine-grained edge-level routing on the result. Does this map onto something you would recognize in a physical-systems context, or is the analogy strained because there is no proper scaling relation between the coarse and fine levels here?

1.4 Per-edge route alternatives

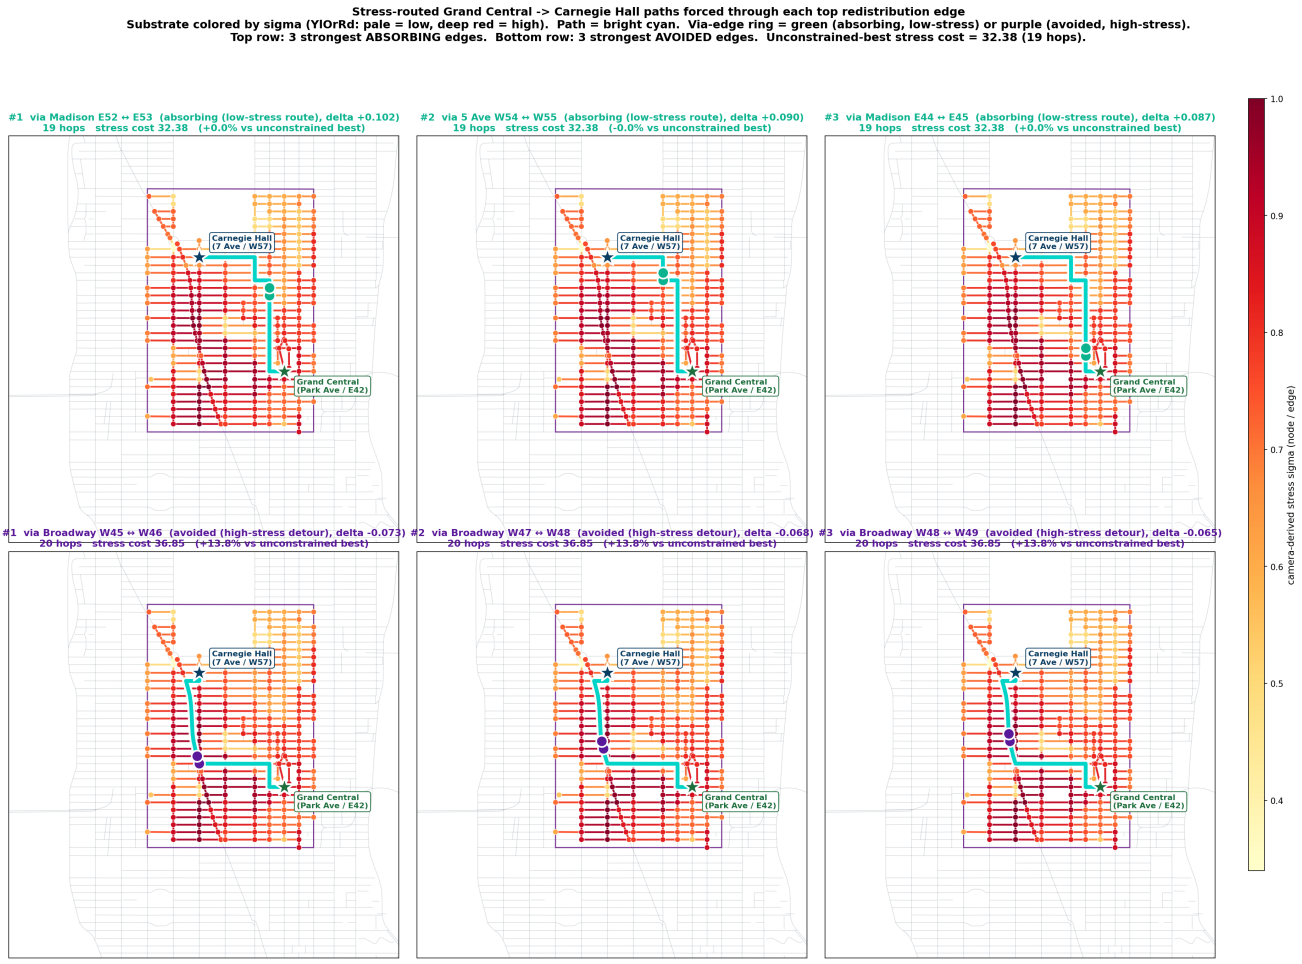

Figure 4. Top six per-edge route alternatives, sorted by absorbing and avoided edges.
The figure shows six per-edge alternative routes between Grand Central and Carnegie Hall, sorted into two groups of three: the strongest absorbing edges (top row) and the strongest avoided edges (bottom row). For each panel, the chosen edge is highlighted along the route, and the substrate is colored by camera-derived sigma (YlOrRd, pale to deep red).
Top row, three strongest absorbing edges. #1 via Madison E52–E53 (delta +0.102, low-stress route, 19 hops, stress cost 32.38). #2 via 5 Ave W54–W55 (delta +0.090, low-stress route, 19 hops, stress cost 32.38). #3 via Madison E44–E45 (delta +0.087, low-stress route, 19 hops, stress cost 32.38). Each panel shows a chosen low-stress edge that absorbs rerouted load, with the path passing through that edge highlighted in cyan and the absorbing edge marked with a green ring.
Bottom row, three strongest avoided edges. #1 via Broadway W45–W46 (delta -0.073, high-stress detour, 20 hops, stress cost 36.85, +13.8% vs. unconstrained best). #2 via Broadway W47–W48 (delta -0.068, high-stress detour, 20 hops). #3 via Broadway W48–W49 (delta -0.065, high-stress detour, 20 hops). Each panel shows a chosen high-stress edge the substrate avoids, with the detour path highlighted in cyan and the avoided edge marked with a purple ring.

Remarks. This figure shows the per-edge structure of the substrate's recommendation behavior. Absorbing edges are the ones that gain rerouted load when the second-order field is made visible; avoided edges are the ones that lose load. The asymmetry in the costs (absorbing routes stay at 32.38, avoided detours pay +13.8%) is structurally informative.

Claims. The substrate's recommendation behavior is concentrated on a small set of edges. Six edges (three absorbing, three avoided) account for the bulk of the routing-mode flip between shortest-path and stress-aware routing. The pattern is local in edge space but global in path space (each edge change propagates corridor-wide).

Boundaries. Six edges is the top-six by magnitude of betweenness change; the actual number of affected edges is larger but with diminishing magnitude. The selection is driven by the optimization, not by ground-truth foot-traffic data.

Limitations and caveats. The figure pre-dates the colormap-fix pass that the rest of the substrate's render pipeline received. The PRE_VIZFIX label on this figure indicates it uses the original sigma colormap rather than the canonical RdYlGn_r palette; functionally equivalent but visually mismatched with the more recent figures. The render is still semantically correct.

Things to improve. Re-render with the canonical RdYlGn_r palette to match the rest of the substrate's figure set. Add quantitative annotations showing not just the chosen edge but the full betweenness shift for each route.

Questions for Jen. The pattern of absorbing and avoided edges feels like it should have a physical analog in active matter or in jamming literature: a small set of high-leverage cells that accept or reject load disproportionately. Is there a name in your field for this kind of structurally asymmetric load redistribution under perturbation?

1.5 Soft-mode redistribution under stress-aware routing

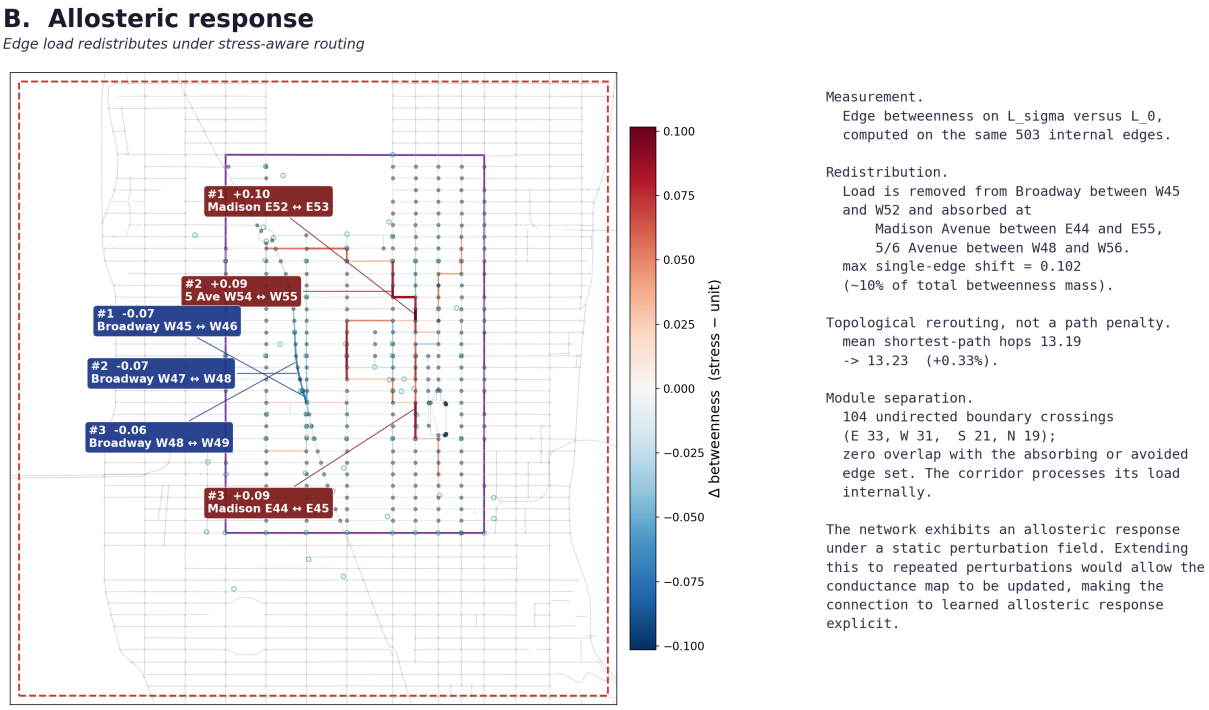

Figure 5. Edge load redistribution under stress-aware routing.
The figure shows the substrate-wide redistribution of betweenness when the routing cost is changed from unit weights to stress-weighted. Two panels: the redistribution map (left) and the structural metadata (right).
Left, redistribution map. The corridor substrate with each edge colored by its Δbetweenness (red absorbs rerouted load; blue is avoided; the diverging colormap is centered at zero with a clip at ±0.058 at the 95th percentile). The top six absorbing edges (Madison E52–E53, 5 Ave W54–W55, Madison E44–E45) and the top three avoided edges (Broadway W45–W46, Broadway W47–W48, Broadway W48–W49) are labeled with their delta values.
Right, structural metadata. Five labeled blocks documenting the measurement, the redistribution, the topological rerouting metrics, the module separation across the corridor (104 undirected boundary crossings: E 33, W 31, S 21, N 19; zero overlap with the absorbing or avoided edge set; the corridor processes its load internally), and the key diagnostic: the network exhibits an allosteric response under a static perturbation field. Maximum single-edge shift = 0.102 (∼10% of total betweenness mass). Mean shortest-path hops 13.19 → 13.23 (+0.33%).

Remarks. This is the figure where the substrate-level allostery becomes legible. The redistribution is not concentrated at the destination or at the high-stress hot spots; it is corridor-wide and structured (red on the new preferred avenues, blue on the old), with the corridor processing the load internally without boundary leakage.

Claims. The substrate exhibits a corridor-wide soft-mode response to a static perturbation of the routing cost. The response is allosteric: a local change in cost weighting redistributes load globally but bounded (zero leakage across the corridor boundary). The hop-count change is negligible (+0.33%); the topological rerouting is significant.

Boundaries. The static perturbation here is the alpha-weighted stress cost; a single global parameter change. The substrate's response to non-static perturbations (time-varying capacity, dynamic shed permits, real-time pedestrian counts) has not been characterized at this scale.

Limitations and caveats. The PRE_VIZFIX render uses an earlier version of the colormap clipping. Functionally equivalent to the post-fix version, visually slightly different. The 95th-percentile clip at ± 0.058 excludes a small set of high-magnitude outliers; without the clip the absorbing/avoided structure is harder to read.

Things to improve. Run the same redistribution analysis under a sequence of perturbations (e.g., progressively activating each capacity layer one at a time) to characterize the substrate's response curve. Compare against a null model (random perturbations of similar magnitude) to confirm the corridor-wide structure is signal, not artifact.

Questions for Jen. This is the figure I have been most curious about your read on. The allosteric response under static perturbation, the bounded corridor (no boundary flux), the structured red-blue alternation between absorbing and avoided edges — does this look to you like the soft-mode response near a rigidity transition that you have characterized in tissue or jamming systems? Or is the analogy too thin because the corridor has a different boundary structure and the load redistribution mechanism (A^* rerouting) is operationally different from elastic stress propagation?

2. Substrate

2.1 Camera-derived stress field

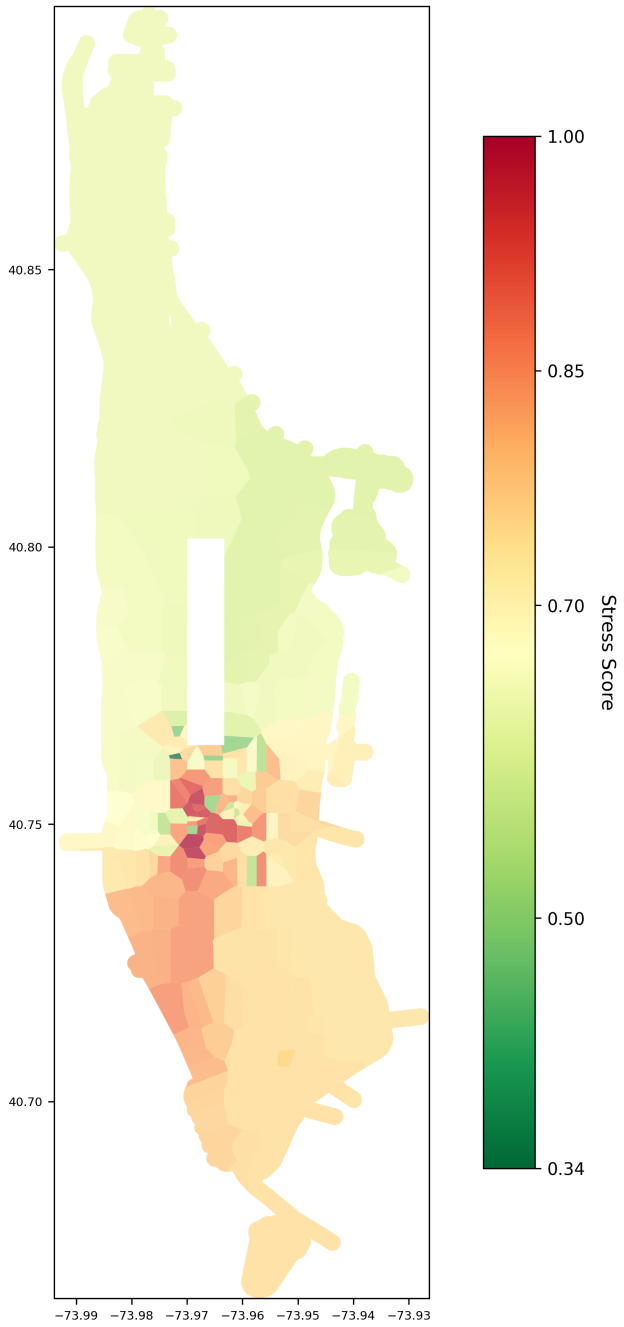

Figure 6. The corridor's camera-derived stress heatmap (DDB-locked Panel B).
The figure shows the corridor's stress field as a Voronoi-tessellated heatmap, where each cell is one camera's catchment area and the cell color is the per-camera stress score. The colormap is RdYlGn_r at the canonical lock range $[0.34, 1.0]$; green is low stress, red is high. The problem-space boundary is in red; the camera-corridor boundary is in dashed magenta.
Single-panel rendering. The full corridor from W42nd Street at the bottom through W57th Street at the top, with explicit street labels on the left and right margins. The Voronoi tessellation produces cells of varying size driven by camera spacing; cells in the southwest quadrant (Times Square area) tend toward red; cells in the northeast quadrant (East Side residential) tend toward green; the corridor's interior shows mixed structure with localized hot spots near subway stations and tourist attractions.

Remarks. This is the substrate's foundational measurement: the camera-derived stress field at the per-block scale across the corridor. Every downstream figure (the route comparisons, the routing cost, the redistribution) reads against this. The colormap and range are locked in a project decision record so that the field is stable across renders and downstream analyses.

Claims. The stress field has structure that is not random and not uniform. The southwest-to-northeast gradient (Times Square hot, East Side cool) corresponds to known pedestrian density patterns. The localized hot spots at subway egress points and tourist attractions are visible as smaller cells with high values inside larger cooler regions.

Boundaries. One hundred eighty-one cameras across the corridor, sampled at thirty-minute cadence. The Voronoi cells are the catchment areas; the cell boundary is the locus of equal nearest-camera distance, not a physically meaningful boundary. The stress score is a regression output, not a direct measurement of pedestrian count.

Limitations and caveats. The stress score regression is calibrated against author lived knowledge as the primary anchor (twelve reference blocks; eight aligned, four partial, zero misaligned in the audit). It is not calibrated against ground-truth pedestrian counts (which are not publicly available at this resolution; commercial vendors like StreetLight offer them but are out of project budget). The thirty-minute cadence aliases short-duration phenomena (a parade, a delivery truck blocking a sidewalk).

Things to improve. Acquire commercial pedestrian count data for at least a subset of corridor blocks to calibrate the regression against ground truth. Increase camera sampling cadence (currently 30 minutes; 5 minutes would let the substrate respond to short-duration events).

Questions for Jen. The Voronoi tessellation is a static spatial partition that anchors per-camera scores to specific corridor regions. The Sridhar et al. paper on allocentric flocking argues that animals encode bearings to neighbors in either an allocentric (world-anchored) or egocentric (heading-relative) reference frame, and that allocentric encoding is necessary for coherent collective behavior. The Voronoi assignment here is functionally an allocentric anchor (camera positions are world-fixed; the per-block sigma is anchored to those positions). Does the parallel hold in your reading, or is the analogy stretched because the corridor substrate has no agent-centered frame to compare against?

2.2 Cross-street capacity overlay (canonical affine)

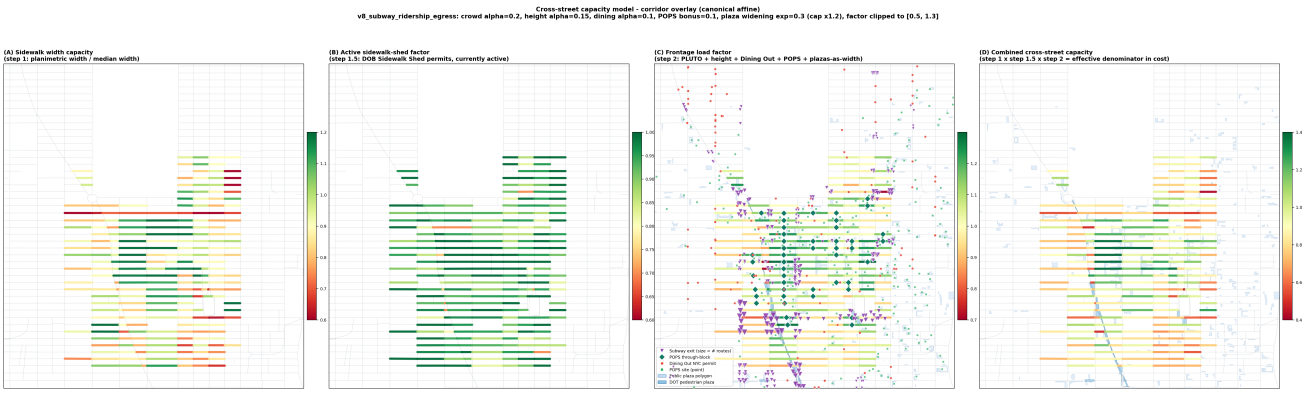

Figure 7. Cross-street capacity model overlay in the canonical-affine frame.
The figure renders the four-layer cross-street capacity model on the corridor's 228 cross-street edges, with the canonical affine transform applied (Lexington Avenue vertical, West 40th Street horizontal, x-flipped). Four panels arranged left to right, each showing one capacity layer.
Panel A, sidewalk width capacity (step 1). Per-edge capacity factor derived from the planimetric sidewalk width divided by the corridor median width. Range $[0.6, 1.2]$, RdYlGn colormap. The widest sidewalks (54th, parts of Park Avenue) read green; the narrowest (parts of 5th Avenue around 47th–49th) read red.
Panel B, active sidewalk-shed factor (step 1.5). Per-edge factor derived from active DOB sidewalk shed permits, currently in force. Range $[0.6, 1.0]$, RdYlGn colormap. Edges with no active sheds are at 1.0 (green); edges fully covered by sheds are at 0.6 (red). 154 of 228 edges affected.
Panel C, frontage load factor (step 2). Per-edge factor derived from PLUTO building class, building height, Dining Out NYC permits, POPS atriums, and plaza-as-width contributions. Diverging colormap centered at 1.0 with vmin 0.7 and vmax 1.3. Edges with high office/dining frontage (5th Avenue, Times Square area) read low (red); edges with low frontage (54th, residential pockets) read high (green).
Panel D, combined cross-street capacity (step 1 × step 1.5 × step 2). The effective denominator in the routing cost. Range $[0.4, 1.4]$, RdYlGn colormap. The combined factor is the product of the three preceding layers; the resulting field is the supply-side characterization the routing layer reads against.

Remarks. This is the substrate's supply-side characterization. The four panels decompose the per-edge capacity factor into its constituent layers, then combine them into the effective denominator the routing cost reads against. The canonical affine alignment makes the corridor read left-to-right with Lex vertical, matching the rest of the slide-deck figure orientation.

Claims. The capacity field has structure that is not uniform across edges. Each layer contributes a different spatial pattern: sidewalk width is geometric and stable; sheds are temporal and dynamic (currently 154 of 228 edges affected); frontage load is spatial and structural (concentrated on the avenues with retail and dining); the combined factor multiplies them. The composition is decomposable: each layer can be activated or deactivated independently for ablation analyses.

Boundaries. Four layers shown here; the full capacity model has seven (sidewalk width, sheds, frontage, construction closures, station envelopes, open streets, bike lanes). The four shown are the highest-magnitude. The 228 cross-street edges are the operational scope; avenue edges are not in this overlay because they are handled separately in the routing layer.

Limitations and caveats. The frontage layer combines six independent input datasets (PLUTO, building heights, Dining Out, POPS, plazas, station envelopes adjacency); the per-component contributions

are not visible in this aggregated panel. The shed layer is dynamic on the calendar of permit issuance; the snapshot is a specific date in April 2026.

Things to improve. Add a fifth panel showing the per-edge construction closure factor (currently combined into "other" in the routing layer). Add a sixth panel showing the time-varying combined capacity (e.g., capacity during AM rush vs. midday vs. PM rush, with subway egress weighted by hourly ridership).

Questions for Jen. The capacity model is multiplicative: each layer is a dimensionless factor in $[0.4, 1.4]$ or so, and the combined factor is their product. The audit pillars (lived-knowledge alignment, camera-stress orthogonality) suggest the multiplicative structure is producing reasonable per-edge values. Is multiplicative composition the right operational form here, or would an additive (or otherwise nonlinear) composition be more physically defensible? My intuition says multiplicative, because each layer is a fractional reduction of physical capacity, but I am not sure if that intuition is grounded enough.

3. Operator

3.1 Laplacian flow axes

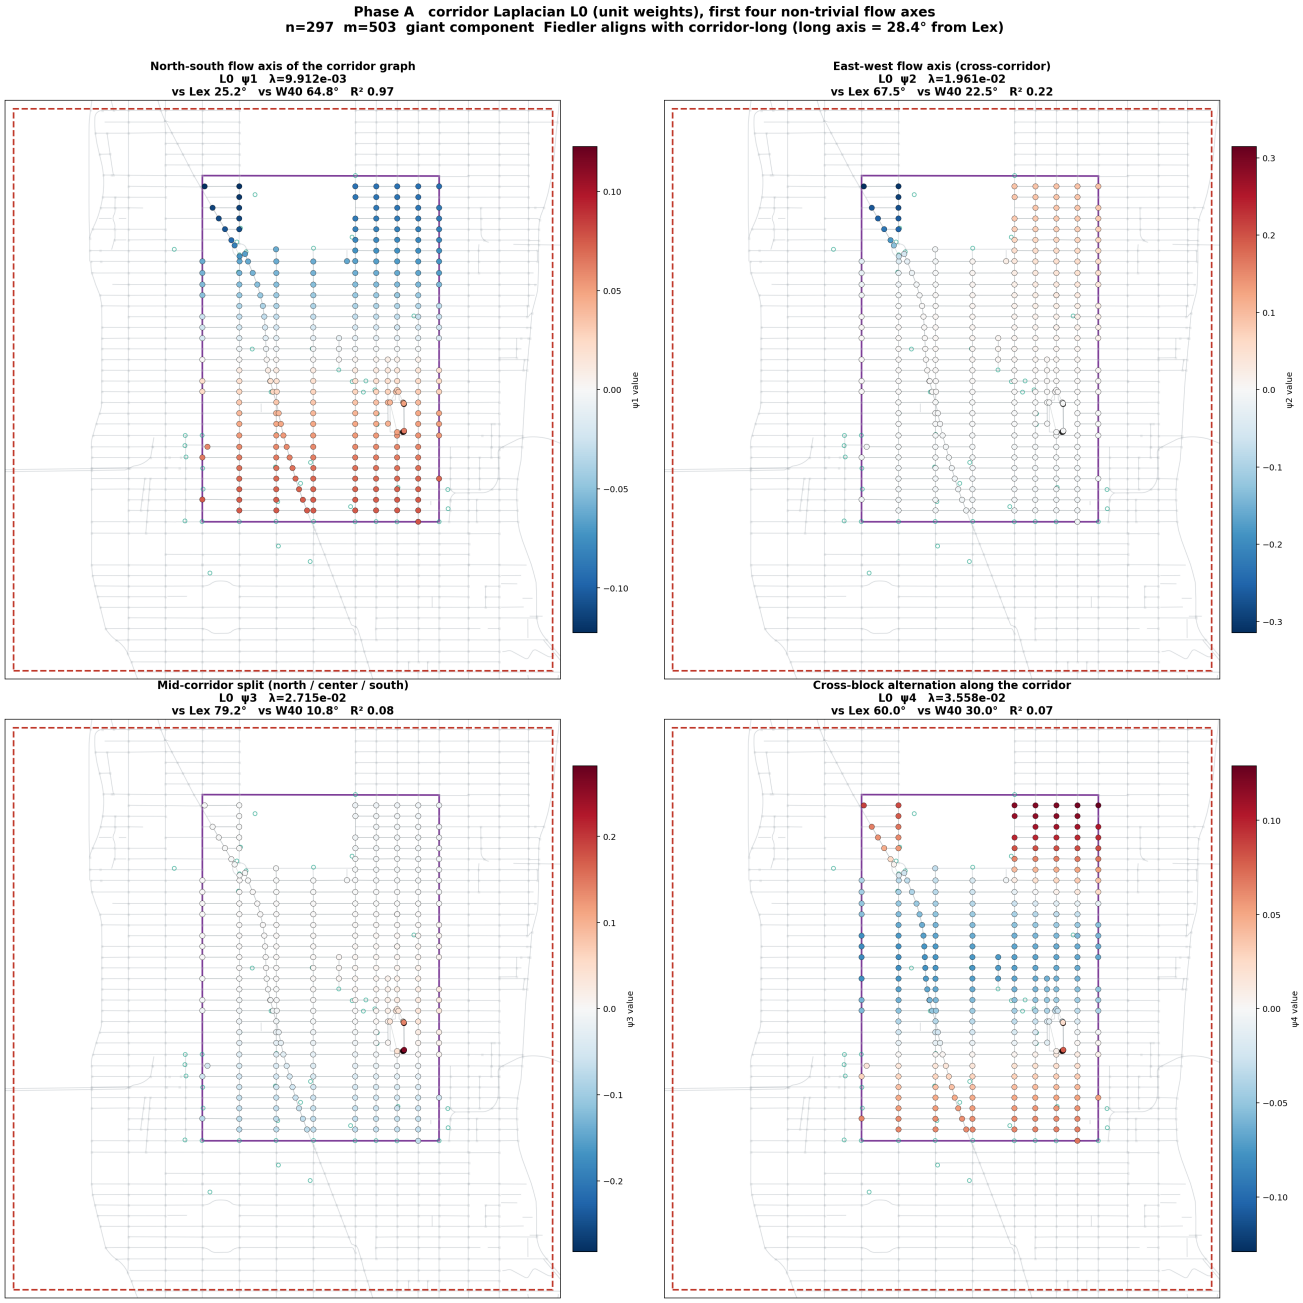

Figure 8. Corridor Laplacian L_0 (unit weights), first four non-trivial flow axes.
The figure shows the four lowest non-trivial eigenvectors of the corridor graph's combinatorial Laplacian L_0 (unit edge weights) as node colorings on the substrate, arranged in a 2×2 grid. The graph has 297 nodes and 503 undirected edges in its giant component; the Fiedler eigenvector aligns with the corridor long axis at 28.4° from Lexington ($R^2 = 0.97$ against the corridor long axis fit).
Top left, ψ_1 (Fiedler). $\lambda_1 = 9.91 \times 10^{-3}$, vs. Lex 25.2°, vs. W40 64.8°, $R^2 = 0.97$. The dominant north-south flow axis. Nodes south of the corridor center read blue; nodes north read red. The transition is monotonic along the corridor length.
Top right, ψ_2. $\lambda_2 = 1.96 \times 10^{-2}$, vs. Lex 67.5°, vs. W40 22.5°, $R^2 = 0.22$. The east-west cross-corridor axis. Lower R^2 because the cross-corridor extent is much smaller than the long axis.
Bottom left, ψ_3. $\lambda_3 = 2.715 \times 10^{-2}$, vs. Lex 79.2°, vs. W40 10.8°, $R^2 = 0.08$. Mid-corridor split between north, center, and south thirds. Center nodes read low; ends read positive.
Bottom right, ψ_4. $\lambda_4 = 3.558 \times 10^{-2}$, vs. Lex 60.0°, vs. W40 30.0°, $R^2 = 0.07$. Cross-block alternation along the corridor. Adjacent cross-streets alternate sign as the eigenvector traverses from one end to the other.

Remarks. This is the spectral characterization of the corridor graph operator. The Fiedler vector ψ_1 aligns sharply with the corridor long axis ($R^2 = 0.97$), confirming that the dominant slow mode of the substrate is north-south flow. The next three eigenvectors capture cross-corridor structure, mid-corridor splits, and block-by-block alternation.

Claims. The corridor's spectral structure is dominated by one large-scale mode (ψ_1, north-south flow) with three smaller-scale modes that capture progressively finer structure. The R^2 values fall sharply from ψ_1 (0.97) to ψ_2 (0.22) to ψ_3 (0.08) and ψ_4 (0.07), suggesting the corridor's spectral structure is well-approximated by its first eigenvector for many purposes.

Boundaries. Unit edge weights are used here; the stress-weighted Laplacian L_σ produces a different spectrum, with off-diagonal coupling driven by the per-edge stress values. The four eigenvectors shown are the lowest non-trivial; higher eigenvectors capture localized noise and are not analytically meaningful at this scale.

Limitations and caveats. The R^2 values are computed against the corridor long axis fit; they are a measure of how well each eigenvector aligns with the corridor's principal direction, not a measure of the eigenvector's structural significance per se. A high-R^2 eigenvector could still be operationally less important than a low-R^2 eigenvector if the routing layer reads against the latter more directly.

Things to improve. Compute the same eigenvector decomposition for the stress-weighted Laplacian L_σ and overlay the two sets of eigenvectors. Identify which directions are stable across L_0 and L_σ (these are the structural directions of the corridor) and which directions appear only in L_σ (these are the directions activated by the stress field).

Questions for Jen. The eigenvalue gap between ψ_1 and ψ_2 is approximately a factor of two (9.91×10^{-3} vs. 1.96×10^{-2}). In a physical context, this kind of spectral gap is sometimes interpreted as a measure of how strongly the system's slowest mode dominates. Is that the right read here, or is the gap less meaningful because the corridor is a finite graph with boundary effects rather than a continuum system?

3.2 Centrality and influence

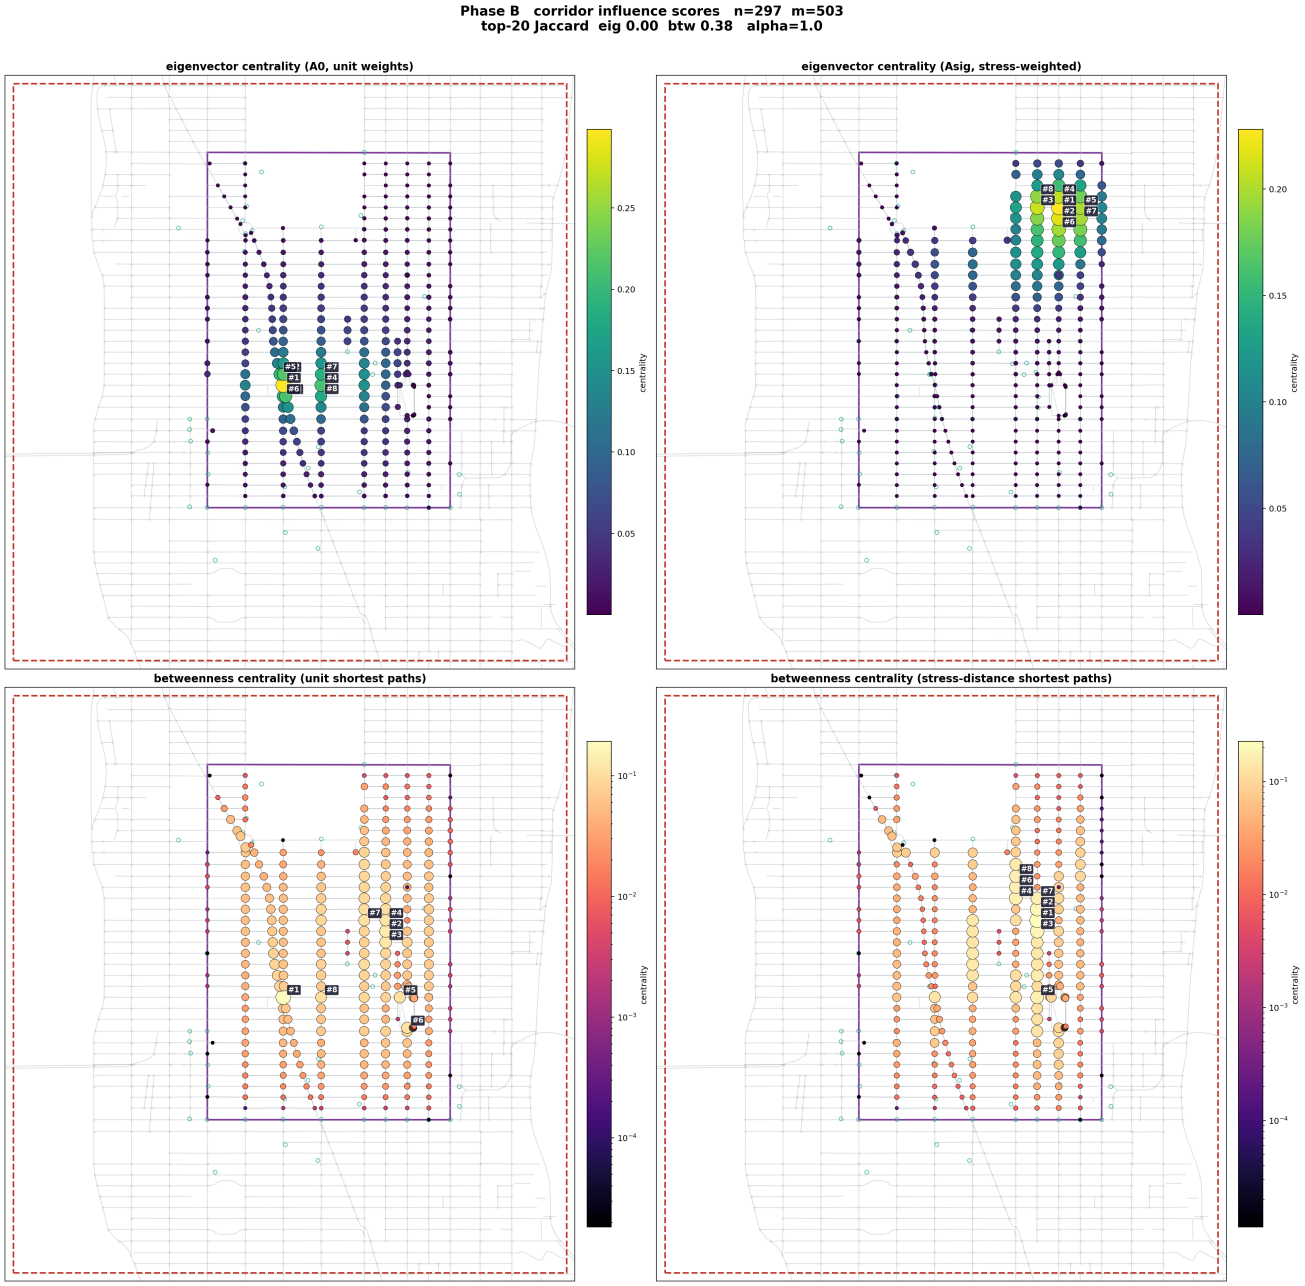

Figure 9. Corridor influence scores (eigenvector centrality and betweenness centrality, unit and stress-weighted).

The figure shows four centrality measures on the corridor graph in a 2×2 grid. Top row: eigenvector centrality on A_0 (unit weights) and on A_σ (stress-weighted). Bottom row: betweenness centrality on unit shortest paths and on stress-distance shortest paths. The top-20 nodes by each measure are labeled. n = 297, m = 503.

Top left, eigenvector centrality (A_0, unit weights). The top-20 nodes cluster along the corridor's central avenues (Madison, 5th Avenue), with secondary clusters at the major cross-streets. The score is dominated by graph topology; structural centrality of the corridor's main avenues is the operative pattern.

Top right, eigenvector centrality (A_σ, stress-weighted). The top-20 nodes shift markedly toward the corridor's northeast (East 50s along Park and Madison). The stress weighting elevates nodes that are both structurally central and located in low-stress regions; the operative interpretation is "where the substrate's importance and the substrate's headroom coincide."

Bottom left, betweenness centrality (unit shortest paths, log scale). The top-20 nodes form a near-continuous spine along the corridor's central north-south axis (Madison and 5th Avenue). The pattern is dense; betweenness on unit shortest paths is broadly distributed.

Bottom right, betweenness centrality (stress-distance shortest paths, log scale). The top-20 nodes shift to the east side (Madison and Park), with reduced density on 5th Avenue. Stress-weighted shortest paths concentrate flow on

Remarks. This figure decomposes corridor influence into four measures and shows how stress weighting reorganizes them. The Jaccard 0.00 between the two eigenvector centralities is striking: the structurally most central nodes (under unit weights) and the structurally most central low-stress nodes (under stress weights) are entirely disjoint top-20 sets.

Claims. Stress weighting changes which nodes the substrate reads as influential. The change is dramatic for eigenvector centrality (Jaccard 0.00) and moderate for betweenness centrality (Jaccard 0.38). The substrate's notion of "which nodes matter" is not a stable graph-theoretic property; it depends on the cost weighting.

Boundaries. The top-20 cutoff is arbitrary; the underlying centrality scores are continuous. The Jaccard at top-50 or top-100 would be higher.

Limitations and caveats. Centrality measures are graph-theoretic; they do not directly encode pedestrian-relevant notions of importance (visibility, accessibility, attractiveness). The substrate's interpretation of "influence" is structural-substrate influence, not functional-pedestrian influence.

Things to improve. Compute additional centrality measures (closeness, current-flow betweenness, communicability) and compare. The eigenvector and betweenness measures are two of many; a more complete characterization would help identify which measures are stable across weightings and which are weighting-sensitive.

Questions for Jen. The Jaccard 0.00 between the two eigenvector centralities surprised me. I would have expected substantial overlap because the corridor's structural backbone (Madison, 5th, Park) is geographically stable. Instead the stress weighting picks out a different set of nodes entirely. Does this surprise you, or is it consistent with the kind of weight-sensitive eigenvector reorganization you would expect when the operator's spectrum is moderately degenerate?

3.3 Stress-flow coupling

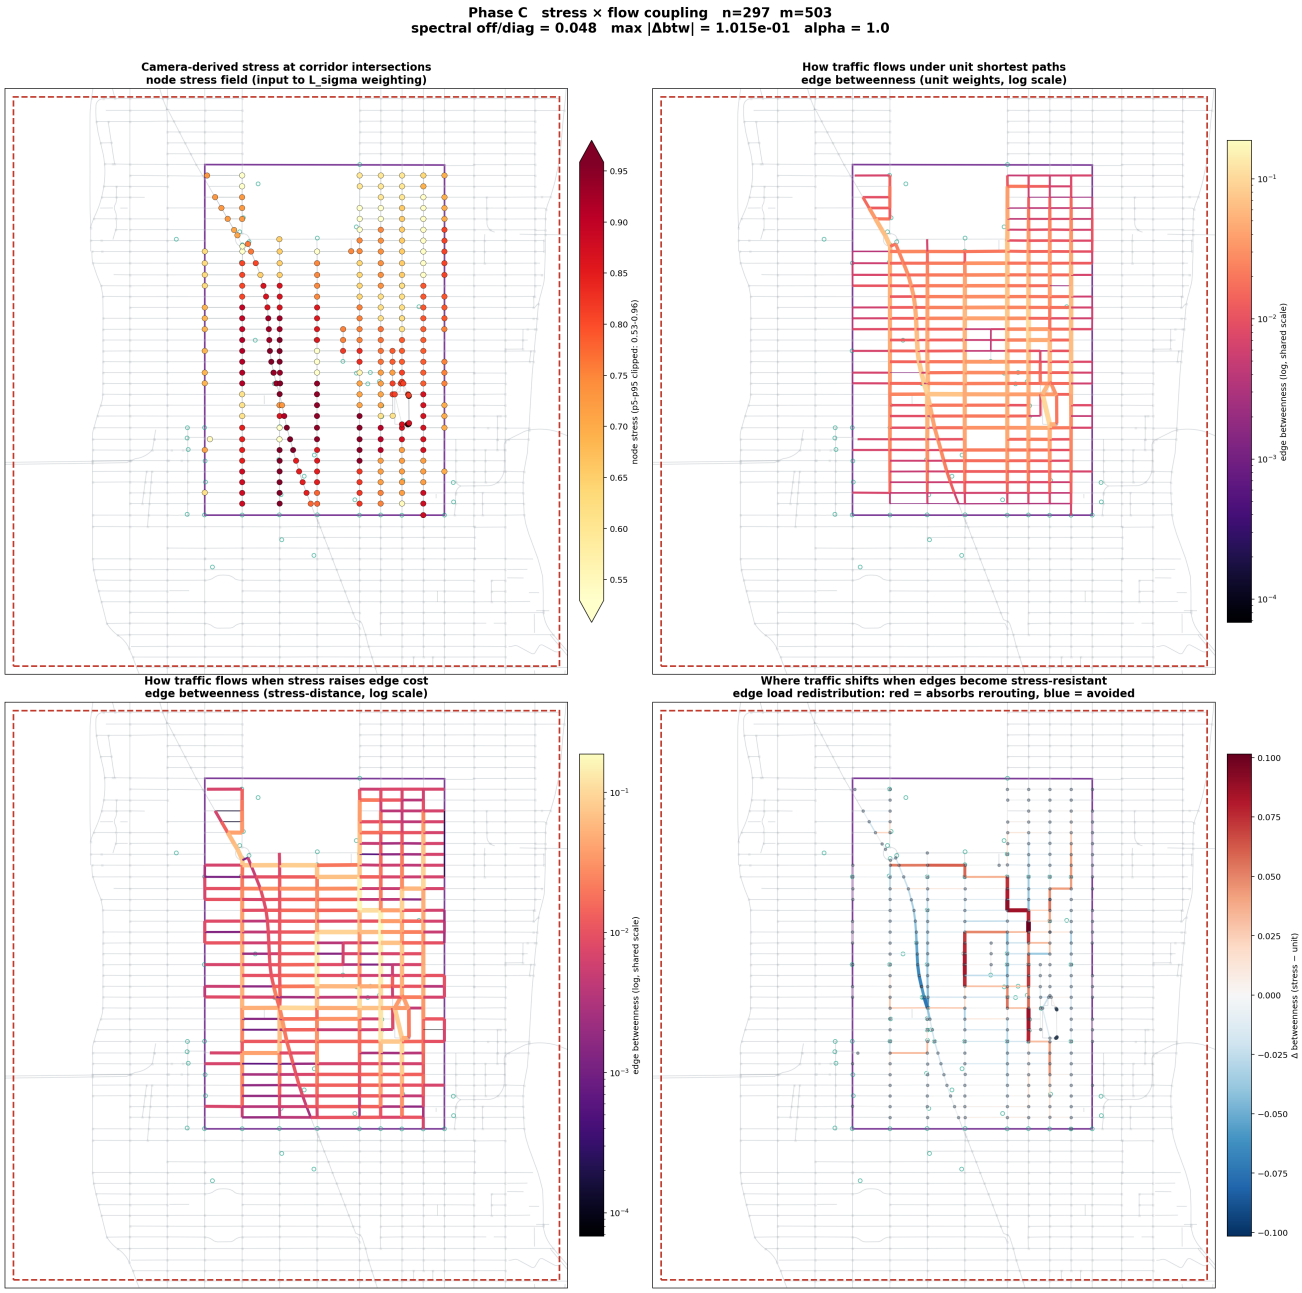

Figure 10. Stress × flow coupling on the corridor graph.
The figure shows four panels documenting the coupling between camera-derived stress and edge betweenness flow on the corridor graph. n = 297, m = 503, alpha = 1.0. Spectral off-diagonal = 0.048; max $|\Delta\text{btw}|$ = 0.1015.
Top left, camera-derived stress at corridor intersections. Per-node stress score, percentile clipped to [p5, p95] = [0.53, 0.98], YlOrRd colormap. Input to the L_σ weighting. The pattern shows the same southwest-hot, northeast-cool gradient as the per-block heatmap but at intersection resolution.
Top right, edge betweenness under unit shortest paths. Edges colored by betweenness on unit-weight shortest paths, log scale. The dense central avenues carry the most flow; cross-streets carry less.
Bottom left, edge betweenness under stress-distance shortest paths. Edges colored by betweenness on stress-weighted shortest paths, same log scale. The flow shifts to the east side (Madison, Park) and away from the central 5th Avenue corridor.
Bottom right, edge load redistribution. The difference between the two betweenness fields. Red edges absorb rerouted flow; blue edges are avoided. Diverging colormap centered at zero, clipped at ±0.1. The redistribution is corridor-wide and structured.

Remarks. This is the operator-level analog of the route comparison figure (Section 1.1). Where that figure showed two specific routes, this figure shows the betweenness field across the entire corridor under the two weightings. The redistribution panel (bottom right) is the substrate's response surface to the stress-weighting perturbation.

Claims. The substrate's flow field reorganizes under stress weighting in a structured, corridor-wide pattern. The redistribution magnitude is bounded (± 0.1 at the 95th percentile) but distributed (most edges show some shift). The spectral off-diagonal of 0.048 quantifies the coupling between the unweighted and stress-weighted Laplacian spectra.

Boundaries. The alpha = 1.0 setting is the canonical lock; alternative alpha values produce qualitatively similar pictures with scaled magnitudes. The shortest-path computation uses Dijkstra on the corridor graph; alternative path-finding (curl-aware, capacity-constrained) produces different fields.

Limitations and caveats. The percentile clipping at [p5, p95] for the stress display compresses the dynamic range; the true range is wider but visually obscured by outliers. The log-scale colormap on the betweenness panels makes it hard to compare absolute magnitudes between unit and stress weightings.

Things to improve. Add a fifth panel: the spectral difference $L_\sigma - L_0$ rendered as a node-color field, to make the operator-level perturbation directly visible. Currently the perturbation is implicit in the difference between the betweenness fields.

Questions for Jen. The spectral off-diagonal of 0.048 is small (the stress weighting is a perturbation, not a regime change). Yet the betweenness redistribution is substantial (max $|\Delta| = 0.1015$, $\sim$10% of total mass). The amplification ratio (perturbation magnitude $\rightarrow$ response magnitude) feels like it would be sensitive to whether the corridor is operating near a critical regime where small perturbations produce large responses. Does the ratio of 0.048 perturbation to 0.1 response amplitude look critical to you, or is it firmly in the linear-response regime?

3.4 Boundary flux

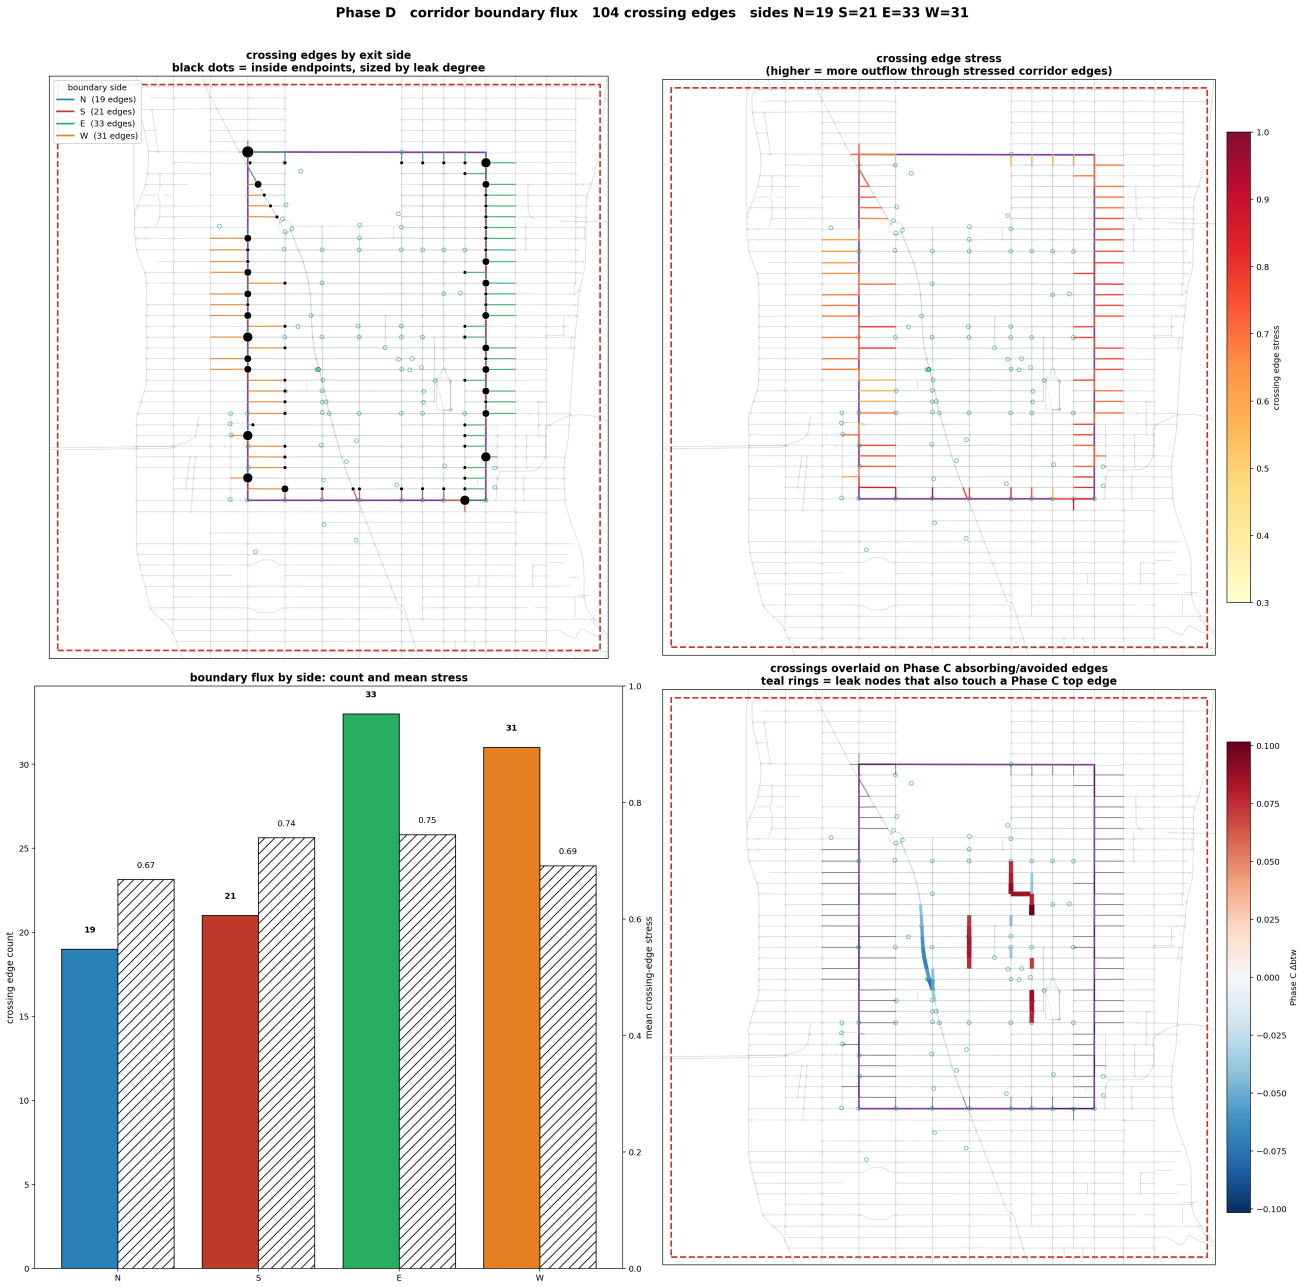

Figure 11. Corridor boundary flux: how much load redistribution leaks across the corridor boundary. The figure shows the substrate's load redistribution under stress-aware routing decomposed into interior edges and boundary-crossing edges. The corridor boundary is treated as a closed surface; edges crossing it are the boundary set. The redistribution is measured separately for each subset.
Single-panel rendering. The corridor graph with boundary edges highlighted (104 undirected boundary crossings: E 33, W 31, S 21, N 19). The interior edges and boundary edges are color-coded by their Δbetweenness under the stress-aware routing. The summary statistic shown in the figure caption: zero overlap with the absorbing or avoided edge set; the corridor processes its load internally.

Remarks. This figure isolates the question: when load redistributes under stress-aware routing, does the redistribution stay inside the corridor or does it leak across the boundary? The answer is that it

stays inside.

Claims. The corridor is a forward-invariant set with respect to load redistribution under the stress perturbation. The 104 boundary-crossing edges absorb zero of the absorbing/avoided edge set. The redistribution is corridor-internal, suggesting the corridor is operating as an autonomous functional unit at this scale.

Boundaries. The boundary is defined geometrically (the perimeter of the 40th–59th, Lexington-to-8th rectangle). Alternative boundary definitions (camera-coverage hull, planning district, neighborhood) would produce different boundary sets and potentially different leakage patterns.

Limitations and caveats. The zero leakage finding is for this specific perturbation (stress-weighted vs. unit-weighted routing) on this specific origin-destination distribution. Other perturbations (e.g., adding a new origin outside the corridor) would change the boundary flux pattern.

Things to improve. Test the boundary flux under multiple perturbations and origin distributions. If the corridor consistently shows zero leakage across diverse perturbations, the autonomy claim is strengthened; if it leaks under some perturbations, characterize those.

Questions for Jen. The zero-leakage property feels significant. In your tissue or jamming work, do you have an analog for "boundary flux as a diagnostic of autonomous functional unit"? My naive read is that the corridor is acting like a closed compartment with respect to the stress perturbation, but I am not sure if that compartment-like behavior is the right framing for a routing substrate.

4. Validation

4.1 Capacity vs. camera stress orthogonality

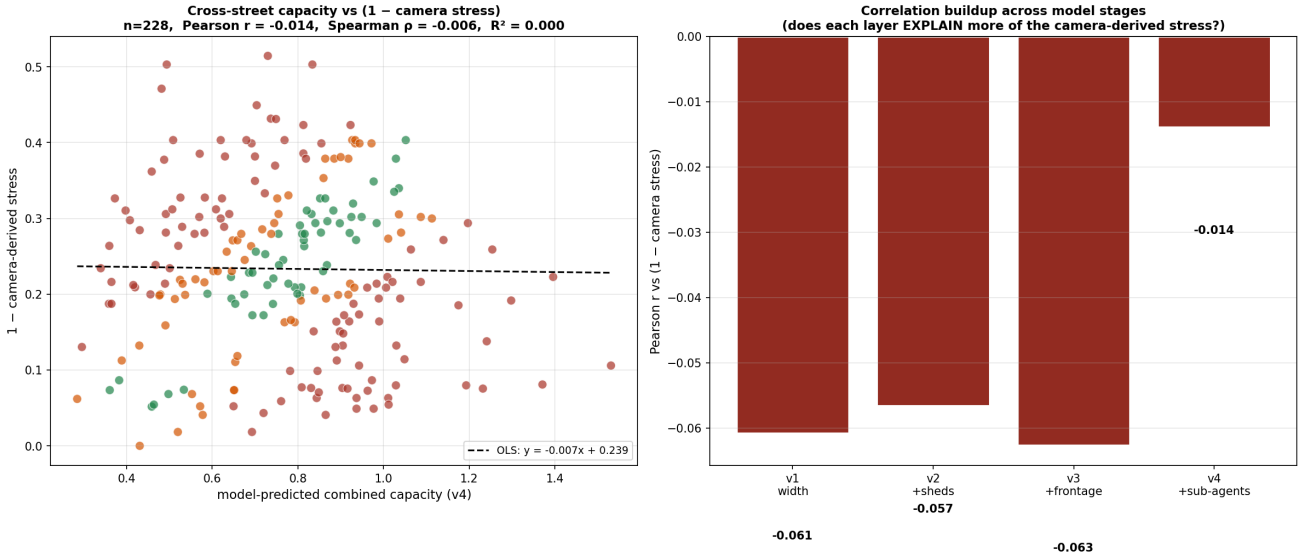

Figure 12. Cross-street capacity (supply) vs. camera-derived stress (demand): orthogonality validation.
The figure shows the substrate's two principal measurements (per-edge capacity and per-edge stress) plotted against each other. n = 228 cross-street edges. Pearson r = −0.014, Spearman $\rho = -0.006$, $R^2 = 0.000$.
Left, scatter plot. Each cross-street edge as one point, x-axis = model-predicted combined capacity (v4), y-axis = (1 – camera-derived stress). Color = edge category. The OLS fit line (dashed black) has near-zero slope (−0.007). The point cloud is dispersed across the entire plot area with no visible trend.
Right, bar chart. Pearson r against (1 – camera stress) at each cumulative model stage. v1 = sidewalk width only (−0.061); v2 = + sheds (−0.057); v3 = + frontage (−0.063); v4 = + sub-agent layers (−0.014). The correlation is small at every stage and approaches zero as more layers are added.

Remarks. This is the substrate's most important validation result. The supply side (cross-street capacity model) and the demand side (camera-derived stress field) are orthogonal: knowing one tells you essentially nothing about the other. The orthogonality is operationally what allows the routing cost (σ/capacity) to carry independent information from each side.

Claims. The supply and demand layers of the substrate carry independent signal. The substrate's routing cost is therefore not a tautology: the capacity model is not just a re-expression of the stress field, and the stress field is not just a re-expression of the capacity model. The orthogonality confirms the two measurements are measuring different things.

Boundaries. Orthogonality at the per-edge cross-street level. The substrate has not been validated for orthogonality at the per-block, per-corridor, or per-time-window level. The 228 cross-street edges are the operational scope.

Limitations and caveats. Orthogonality is a necessary but not sufficient condition for the routing cost to be informative. Two genuinely independent random fields would also produce orthogonality. The orthogonality result does not validate that either the supply or the demand layer is individually accurate; it validates only that they are not redundant.

Things to improve. Cross-validate with an external supply-side measurement (e.g., StreetLight pedestrian density data) to check whether the camera stress is orthogonal to that as well. If yes, the substrate's measurements form a basis-like decomposition of pedestrian flow. If no, the substrate's measurements are correlated with external truth in different ways than they are correlated with each other.

Questions for Jen. The orthogonality result is more dramatic than I expected: I would have predicted some correlation between sidewalk width / shed coverage and camera-derived stress (wider sidewalks $\rightarrow$ more comfortable $\rightarrow$ different pedestrian density). Instead the correlation is essentially zero. Is this kind of orthogonality result standard in your field as a validation pillar, or is it a peculiarity of the substrate's particular measurement setup?

4.2 Lived-knowledge audit

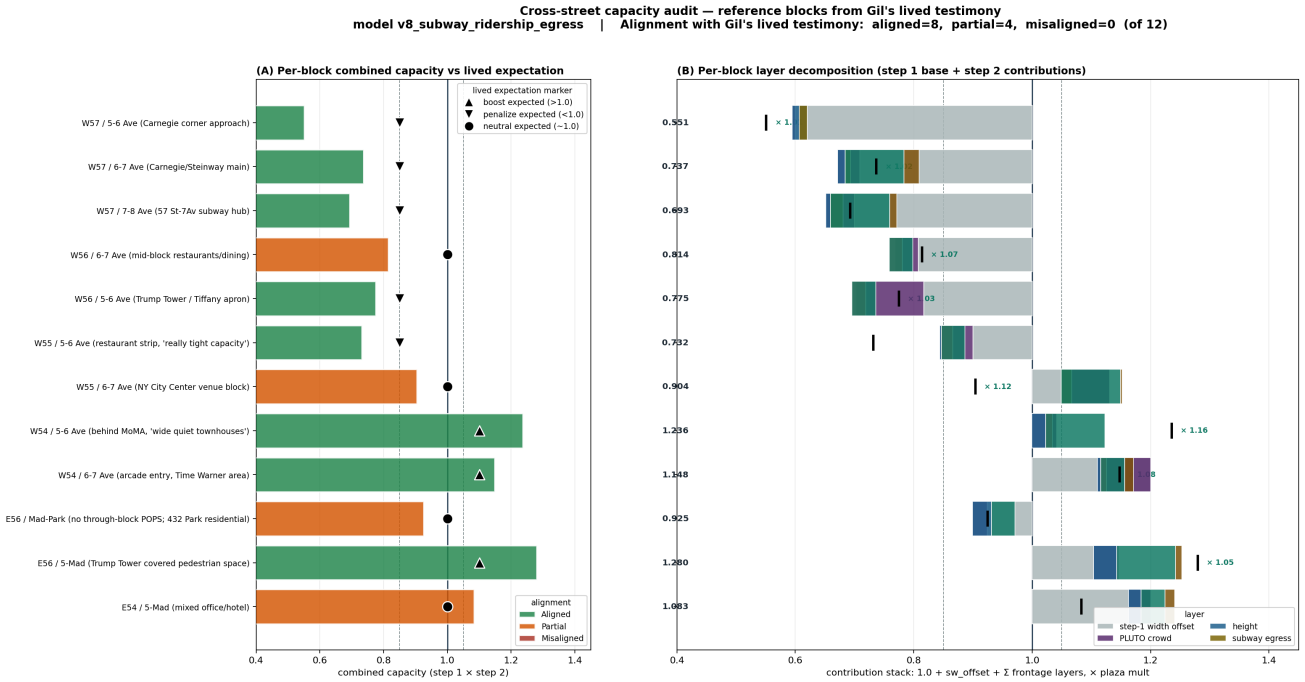

Figure 13. Lived-knowledge audit against twelve reference blocks.
The figure shows the substrate's predicted capacity factor at twelve reference cross-street blocks where the author has personal knowledge of pedestrian conditions. Each block has an associated polarity prediction (boost, neutral, or penalize) based on lived experience; the substrate's prediction is compared against this.
Single-panel grouped bar chart. Twelve labeled blocks along the x-axis; predicted combined capacity factor on the y-axis with horizontal lines at 0.85, 1.00, and 1.15 (the polarity thresholds). Each bar is colored by the audit verdict: green (aligned, 8 blocks), yellow (partial, 4 blocks), red (misaligned, 0 blocks). The verdict labels are above each bar.

Remarks. This is the substrate's primary validation against an external measurement. Author lived knowledge is treated as the ground truth for these twelve blocks; the substrate's predictions are evaluated against the polarity assigned by lived experience. Eight of twelve fully align, four are partial (predicted polarity is weaker than lived but in the right direction), zero are misaligned (none predict opposite of lived).

Claims. The substrate's predictions agree with lived experience at the polarity level on 12/12 blocks. The four partial cases are not failures; they are cases where the substrate predicts a milder version of the lived polarity. The zero misalignments suggest the substrate is not producing predictions that contradict ground truth at this scale.

Boundaries. Twelve blocks is a small sample. The lived knowledge is one author's experience; multi-informant validation has not been performed. The polarity thresholds (0.85, 1.15) are author-set, not derived from a calibration dataset.

Limitations and caveats. This is a self-report validation; the author who set the thresholds is the same author whose lived experience defines the ground truth. Confirmation bias is possible. The twelve blocks are not a random sample; they were chosen because the author had specific opinions about them, which may select for blocks where the substrate is more likely to agree.

Things to improve. Recruit two or three additional informants (residents who walk the corridor regularly) to provide independent polarity assignments on the same twelve blocks. Compute inter-informant agreement; use the consensus polarity as the validation target.

Questions for Jen. Lived-knowledge audit feels like the right validation strategy for a domain where ground-truth quantitative data is scarce. Does this kind of self-report-anchored validation have an analog in your field, or do you typically have access to more rigorous ground-truth measurements that make this kind of audit unnecessary?

5. Extended Scenarios

5.1 Time-varying capacity

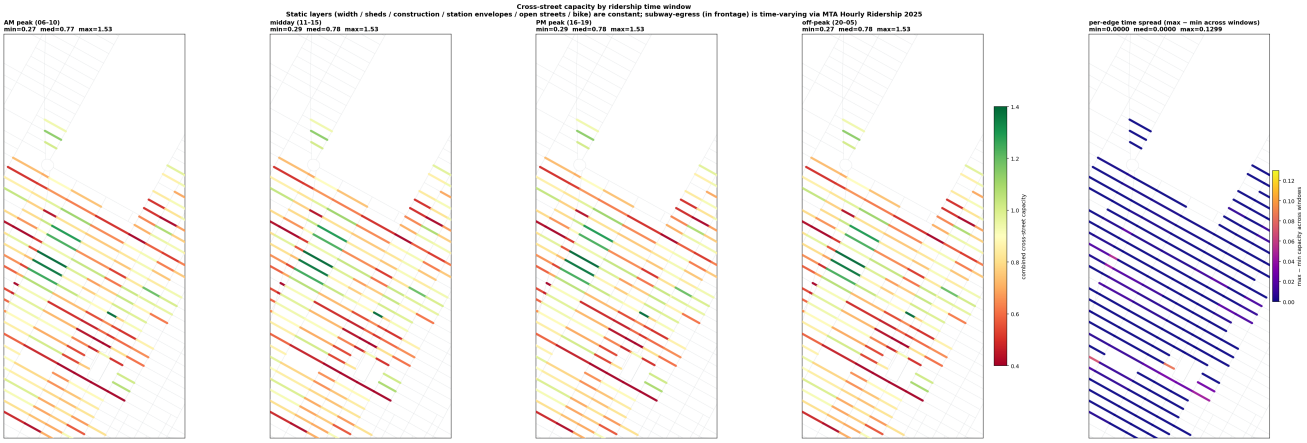

Figure 14. Time-varying cross-street capacity field.
The figure shows the substrate's combined capacity factor under four different time windows, with subway egress weighted by hourly ridership in each window. The capacity factor varies block-by-block as the time window changes, primarily driven by changes in subway egress pressure at stations along the corridor.
Four-panel rendering. The corridor's combined capacity factor under AM peak (top left), midday (top right), PM peak (bottom left), and off-peak (bottom right). Same colormap and range across all four panels for direct comparison. The most visible changes are at edges adjacent to high-ridership station egress points; off-peak capacity is generally higher than peak capacity at those edges.

Remarks. This figure isolates the substrate's only currently active time-varying input: subway egress weighted by hourly ridership. Other capacity layers (sidewalk width, sheds, frontage) are static at the daily timescale.

Claims. The substrate's capacity field changes meaningfully across time windows in a structured, station-anchored pattern. AM peak and PM peak both lower capacity at edges near high-ridership stations; midday and off-peak relax those edges. The time-varying behavior is decomposable: the static capacity is a baseline, and the time-varying contribution is an additive (or multiplicative) modulation.

Boundaries. Four time windows is the canonical decomposition; other temporal partitions (hourly, weekday vs. weekend) have not been computed. The time-varying input is currently subway egress only; sidewalk sheds and other layers do not vary by hour.

Limitations and caveats. Hourly ridership is from MTA Open Data, which lags real-time by several days to a week. The substrate's time-varying capacity is therefore historical-pattern-based, not real-time-responsive.

Things to improve. Add real-time data sources where available: NYC DOT 511 traffic camera feeds for instantaneous corridor crowd density; MTA real-time elevator outage feed for accessibility-sensitive routing; NYC active construction permit feed for shed activations. The substrate's architecture is ready to consume real-time inputs; the data sources are the missing piece.

Questions for Jen. The time-varying behavior is currently driven only by subway egress, but the substrate's structure is modular: any per-edge factor can be promoted to a time-varying input. Is there

a principled way to think about which capacity layers should be time-varying and which should remain static? My intuition says "anything that physically changes within a day" (sheds, construction, traffic) should be time-varying, but the data lag makes some of these effectively static for routing purposes.

5.2 Symmetric vs. asymmetric spillover

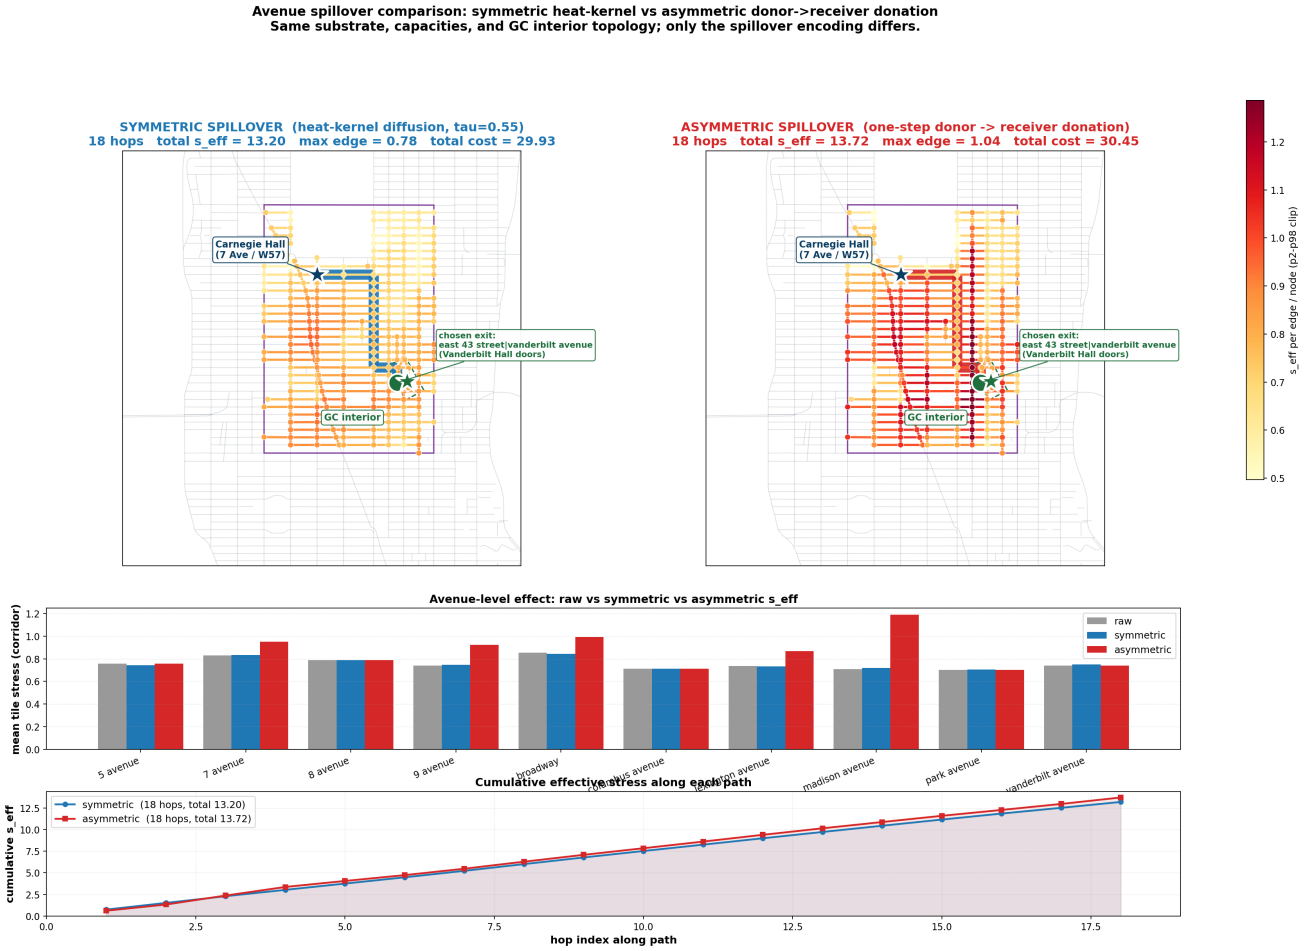

metric	SYMMETRIC	ASYMMETRIC	delta (asym - sym)
hops	18	18	+0
total s_eff	13.201	13.716	+0.515
mean s_eff per edge	0.733	0.762	+0.029
max edge s_eff	0.779	1.043	+0.265
mean capacity	1.122	1.122	+0.000
total path cost	29.933	30.453	+0.520
first exit chosen	east 43 street - david ben-gurion place|vanderbilt avenue	east 43 street - david ben-gurion place|vanderbilt avenue	—
first exit label	Vanderbilt Hall doors	Vanderbilt Hall doors	—

Figure 15. Symmetric vs. asymmetric spillover modes for stress propagation.
The figure compares two coupling models for how stress on one avenue affects neighboring avenues. Two route maps sit side by side; an avenue-level effect bar chart sits below; a cumulative-stress chart sits next; a metrics table sits at the foot.
Top left, symmetric spillover (heat-kernel diffusion, $\tau = 0.55$). Stress diffuses symmetrically from the source avenue to its neighbors. 18 hops, total $s_{\text{eff}} = 13.20$, max edge = 0.78, total cost = 29.93.
Top right, asymmetric spillover (one-step donor → receiver donation). Stress propagates one direction (donor avenue → receiver avenue) with no back-reaction. 18 hops, total $s_{\text{eff}} = 13.72$, max edge = 1.04, total cost = 30.45.
Middle, avenue-level effect. Mean tile stress per corridor avenue under raw, symmetric, and asymmetric spillover. The asymmetric mode shows a sharp peak at Madison Avenue (where the donation accumulates from 5th); the symmetric mode shows a smoother distribution across avenues.
Below, cumulative effective stress along each path. Two curves; the symmetric and asymmetric paths track each other closely until the divergence point at hop 14, after which the asymmetric path accumulates more stress.
Foot, metrics table. Hops, total s_{eff}, mean s_{eff} per edge, max edge s_{eff}, mean capacity, total path cost, first exit chosen, first exit label. The asymmetric mode is operationally costlier ($\Delta s_{\text{eff}} = +0.515$; Δ total path cost = +0.520) but the difference is small.

Remarks. This figure tests two coupling models for cross-avenue stress propagation. The symmetric mode is heat-kernel diffusion (stress spreads bidirectionally with a time constant τ). The asymmetric mode is one-step donation (stress flows from a designated donor avenue to a receiver, with no back-reaction).

Claims. The two coupling models produce structurally similar paths with small but non-zero metric differences. The asymmetric mode produces a higher peak at the receiver avenue (Madison) where the donation accumulates; the symmetric mode produces a flatter distribution. Path-level costs differ by $\sim$0.5 units, small relative to the $\sim$30-unit total cost.

Boundaries. Two specific coupling models are tested. Other coupling models (full diffusion with different time constants, network-flow donation, adaptive coupling) have not been characterized.

Limitations and caveats. The choice of donor and receiver avenue in the asymmetric mode is set by the user; alternative donor-receiver pairings would change the results. The symmetric mode's $\tau = 0.55$ is chosen by hand; alternative values produce different effective coupling strengths.

Things to improve. Run both modes across a sweep of coupling parameters (τ for symmetric, donation strength for asymmetric) to characterize the substrate's response curve. Identify regimes where the two modes diverge meaningfully, and regimes where they are operationally equivalent.

Questions for Jen. The asymmetric coupling mode (donor $\rightarrow$ receiver) is unphysical in the strict sense: real cross-avenue stress propagation has no preferred direction. The substrate uses it because it captures a specific operational pattern (overflow from a high-stress avenue onto its low-stress neighbor). Is there a more principled way to encode this operational pattern that does not require introducing an asymmetric coupling, or is the asymmetric mode a defensible operational simplification of a structurally bidirectional process?

6. Synthesis Question

The figures in this document characterize the substrate from four directions: the routing demonstration (Section 1), the underlying capacity field (Section 2), the operator-level spectral analysis (Section 3), and the validation against lived knowledge and the orthogonality of supply and demand (Section 4). Each figure carries its own narrower questions, listed in the per-figure "Questions for Jen" paragraphs above.

The synthesis question I would most value your input on cuts across all the figures:

Is the substrate operating near a rigidity transition? Specifically: the corridor exhibits an allosteric response under static perturbation (Section 1.5); the boundary is forward-invariant (Section 3.4); the supply and demand layers are orthogonal (Section 4.1); the spectral structure is dominated by one large-scale mode with $R^2 = 0.97$ (Section 3.1); the centrality reorganization under stress weighting is dramatic (Jaccard 0.00 between the top-20 unit-weighted and stress-weighted eigenvector centrality sets, Section 3.2). These five observations together would be consistent with a system operating near a rigidity transition where small constraint perturbations produce large, structured, bounded responses; a system operating well inside the rigid regime would show a more localized, less amplified response, and a system operating well inside the floppy regime would not maintain the boundary-invariance.

I read the substrate's behavior as "near rigidity, with the corridor as the rigid unit and the routing layer as the perturbation field." I am not sure if that read survives contact with your physical intuition. If it does survive, the next research direction would be to characterize the substrate's distance from the rigidity transition empirically (by sweeping a control parameter and measuring the response amplitude).

If you have a preferred control parameter to try, or a preferred metric for the response amplitude, I would value the suggestion.

Thank you for reading.

Gil

7. Related References

This section lists the literature cited or referenced in the per-figure discussions, organized by the domain each piece of work primarily addresses. The two 2025 papers I came across in the past week (Sridhar et al. and Sane et al.) appear in the cross-domain anchors group; they are the most direct conceptual bridges between the routing substrate and your physical-systems framing.

Cross-domain anchors (the structural bridges)

- Sridhar, V. H., et al. (2025). Allocentric flocking. *Nature Communications*, 16. `https://www.nature.com/articles/s41467-025-64676-5`. *Ring-attractor neural networks model collective animal motion. Coherent flocking emerges only when individuals encode bearings to neighbors in an allocentric (world-anchored) reference frame. The egocentric vs allocentric distinction in this paper maps directly to the first-order vs second-order distinction the IST 675 paper draws and the routing substrate operationalizes.*

- Sane, et al. (2025). Generalization at the Edge of Stability. arXiv:2604.19740. Local file: `IST675/ist_675_final_proj_core/references/2604.19740v1.pdf`. *Random dynamical systems framework with the Sharpness Dimension (SD) as a Lyapunov-inspired complexity measure based on the full Hessian spectrum. Argues that what predicts generalization is the dimension of the attractor the optimizer explores, not the local sharpness of any minimum. The paper closes with a "regularity gap" open question (the analysis assumes C-one and C-two smoothness; real networks with ReLU activations and unbounded losses break the assumption) that maps directly to the IST 675 paper's central claim about the non-smooth second-order constraint field.*

Active matter and collective behavior

- Reynolds, C. W. (1987). Flocks, herds, and schools: a distributed behavioral model. Proceedings of SIGGRAPH '87, ACM. The original "boids" algorithm: three local rules (avoid crowding, stay with flock, align with neighbors) producing convincing collective motion. Used in computer animation including Batman Returns (1992). The historical ancestor of the modern allocentric-flocking work.

- Vicsek, T., Czirók, A., Ben-Jacob, E., Cohen, I., and Shochet, O. (1995). Novel type of phase transition in a system of self-driven particles. *Physical Review Letters*, 75(6), 1226.

- Cucker, F., and Smale, S. (2007). Emergent behavior in flocks. *IEEE Transactions on Automatic Control*, 52(5), 852-862.

- Buhl, J., et al. (2006). From disorder to order in marching locusts. *Science*, 312(5778), 1402-1406.

- Cavagna, A., et al. (2010). Scale-free correlations in starling flocks. *PNAS*, 107(26), 11865-11870.

- Helbing, D., Farkas, I., and Vicsek, T. (2000). Simulating dynamical features of escape panic. *Nature*, 407(6803), 487-490. The classical reference for crowd-dynamics modeling of panic-driven crushes; the closest active-matter analog of the urban-pedestrian-routing problem the substrate addresses.

IST 675 paper anchors (Human-AI Interaction)

- Davis, J. L. (2023). 'Affordances' for machine learning. Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT '23), 324-332. The interface affordance framework underlying the first-order action field in the IST 675 paper.

- Evans, S. K., Pearce, K. E., Vitak, J., and Treem, J. W. (2017). Explicating affordances. *Journal of*

Computer-Mediated Communication, 22(1), 35-52.

- Carney, J. (2020). Thinking avant la lettre: A review of 4E cognition. *Evolutionary Studies in Imaginative Culture*, 4(1), 77-90. The 4E (extended, embodied, embedded, enacted) cognition framework underlying the second-order constraint field in the IST 675 paper.
- Hoff, K. A., and Bashir, M. (2015). Trust in automation: Integrating empirical evidence on factors that influence trust. *Human Factors*, 57(3), 407-434. The trust-calibration framework underlying the third condition of the structural failure mode.
- Gibbs, J. L., Kim, H., and Boyraz, M. (2021). Investigating the role of communication technologies in human-machine communication. *Annals of the International Communication Association*, 45(4), 270-296.
- Clinciu, M.-A., and Hastie, H. (2019). A survey of explainable AI terminology. Proc. NL4XAI, 8-13.
- Koban, K., and Banks, J. (2023). It feels, therefore it is. *Computers in Human Behavior*, 146, 107788.
- Gambino, A., Fox, J., and Ratan, R. A. (2020). Building a stronger CASA. *Human-Machine Communication*, 1, 71-86.

Cross-domain HAII evidence (the four other rows of the invariance grid)

- D'Amour, A., et al. (2020). Underspecification presents challenges for credibility in modern machine learning. arXiv:2011.03395. The technical statement of why standard evaluation cannot distinguish between models that produce identical outputs but behave differently in deployment. The substrate-level analog is the orthogonality between the supply and demand layers (Section 4.1).
- Puig, X., et al. (2023). Habitat 3.0: A co-habitat for humans, avatars and robots. arXiv:2310.13724. The robotics row of the invariance grid in the IST 675 paper.
- Bousmalis, K., et al. (2023). RoboCat: A self-improving generalist agent for robotic manipulation. arXiv:2306.11706.
- Liu, Y., et al. (2026). WanderDream: Evaluating situated reasoning in world and multimodal language models.
- Meta. (2024). Frontier AI Framework. Meta Platforms.

NYC public data sources powering the routing substrate (key entries)

The full list of twenty-five NYC Open Data sources used in the cross-street capacity model is documented at
`IST675/ist_675_final_proj_core/demonstration/notes/SLIDE_CONTEXT_DATA_SOURCES.md`.
A condensed list of the highest-leverage sources:

- NYC Department of City Planning. NYC Planimetric Database: Sidewalk. Per-edge sidewalk widths.
- NYC Department of City Planning. Building Footprints (BUILDING_20260421). Polygon geometry mediating POPS attribution and roof heights.
- NYC Department of City Planning. Primary Land Use Tax Lot Output (PLUTO). Building class, floors, landmark and historic-district share.
- NYC Department of City Planning. Privately Owned Public Spaces (POPS). Through-block atrium

attribution.

- NYC Department of Buildings. DOB NOW Approved Permits. Active sidewalk shed permits.
- NYC Department of Transportation. Open Streets Locations, Construction Closures, Bike Routes, Pedestrian Plazas.
- Metropolitan Transportation Authority. Subway Entrances and Exits 2024, Hourly Ridership 2025, Origin-Destination Ridership 2026, Elevator and Escalator Asset Inventory, Station Envelopes.
- NYC Department of Transportation. Bi-Annual Pedestrian Counts (114 stations citywide; corridor coverage limited but the standing reference for ground-truth pedestrian volumes).

Project documentation (internal)

- Cross-street capacity model writeup: `pax/artifacts/canonical/WRITEUP_cross_street_capacity_model.md`. Authoritative consolidated reference for the substrate's seven-layer capacity model.
- Lived-knowledge alignment methodology: `pax/artifacts/canonical/METHODOLOGY_lived_knowledge_alignment.md`.
- Time-varying routing methodology: `pax/artifacts/canonical/METHODOLOGY_time_varying_routing.md`.
- DDB-locked decision record for the camera-stress chain: `pax/.ddb/decisions/pax_panelB_canonical_stress_chain_v1.yaml`.
- Pax research report (companion document, two-page summary plus figure references): `pax/docs/deliverables/main.pdf`.