Transport Network Welfare[1]

Theory, Computation, and Practice with `TransportNetworkWelfare.jl`

Treb Allen Simon Fuchs Woan Foong Wong

Practitioner guide for package API v0.2.0, July 2026

What this guide does

TransportNetworkWelfare.jl ranks small improvements to existing transport links by their effect on economic welfare. It takes directed flows, location shares, and calibrated elasticities as inputs and reports three welfare measures: the Traditional traffic-based statistic, the full realized-cost effect, and the Extended primitive-cost effect. At an efficient benchmark, the traffic share measures the first-order welfare effect even after routes, modes, prices, and economic activity adjust. With congestion or spatial externalities, the full-model effects also account for the welfare consequences of those adjustments. Comparing the realized- and primitive-cost effects measures pass-through. The package reports the model comparisons and numerical checks needed to interpret these differences.

The guide develops this comparison through a reproducible 25-location economic-geography grid with a transit spine. The reported effects are marginal gross benefits, not complete project appraisals. Construction and financing costs, safety, pollution, distributional effects, and transition dynamics must be assessed separately. The worked example runs from a fresh clone without external application data, private credentials, an Overleaf checkout, or Python plotting packages. The U.S. freight and Seattle applications require external data. The Sioux Falls and Westeros builders download and verify their source files rather than storing them in the repository.

AI-use disclosure. Generative AI tools assisted with portions of code and test development, debugging, documentation and prose editing, mathematical checks, and visualization. The authors are responsible for the final material and any remaining errors.

[1] Please cite the paper for the method and the guide and tagged software release for the documentation and code: Treb Allen, Simon Fuchs, and Woan Foong Wong (2026), *From Social Savings to Spatial Equilibrium: Evaluating Transportation Improvements in Quantitative Spatial Models*, manuscript; Treb Allen, Simon Fuchs, and Woan Foong Wong (2026), *Transport Network Welfare: Theory, Computation, and Practice with TransportNetworkWelfare.jl*, practitioner guide for version 0.2.0; and Treb Allen, Simon Fuchs, and Woan Foong Wong (2026), *TransportNetworkWelfare.jl*, version 0.2.0, `https://github.com/sfuchs-de/TransportNetworkWelfare.jl`. Cite the tagged software version or commit used in the analysis. Adaptations should also describe and credit substantive contributions from the code, derivations, example builders, and underlying data.

Contents

Reading guide

The guide has four parts. Part I gives a stand-alone route from data to results. Parts II and III form the detailed guide: Part II develops the economic argument, and Part III turns that argument into data, software, and applications. Part IV collects technical material that is easier to consult separately.

Document structure

Part I: Quick start. Section 1 is a stand-alone operating procedure for a new application. It distinguishes the economic-geography and urban-commuting models, states the input accounting conditions, shows the required file schemas and configuration choices, and gives the commands for validation, analysis, and decomposition. Begin there when the immediate question is whether an existing network dataset can be used with the package.

Part II: Detailed guide—theory and welfare. Section 2 fixes the policy experiment, the three reported welfare measures, and the boundary between model-implied benefits and project appraisal. Section 3 then defines the economic-geography and urban models together with their shared transport equations. Section 4 uses those definitions to establish the traffic result, derive the extended welfare elasticity, and decompose the difference between the two approaches.

Part III: Detailed guide—implementation and applications. Section 5 states the data contract, follows the software calculation, and gives the numerical and reporting checks. Section 6 applies the complete sequence to the public multimodal grid. Section 7 distinguishes the audited U.S. freight results, whose replication requires external data, from the urban and Seattle extensions. Section 8 collects the interpretation and reporting requirements for a result.

Part IV: Reference appendices. The appendices are reference material rather than another stage of the main argument. Appendix A maps notation to code; Appendix B records the computational sequence; Appendix C gives the configuration and schema reference; and Appendix D states the numerical tests and tolerances. Appendices E and F summarize the other package examples and document the sources behind the guide.

Suggested reading paths

Reader's task	Suggested path
Adapt a new dataset	Section 1; then Sections 5 and 6; keep Appendices C and D open while preparing the inputs.
Assess the economics	Sections 2–4; then Appendix A for the equation-to-code map.
Build an urban or public-transit application	The model comparison in Section 1; the urban model and shared transport equations in Section 3; then the urban material in Section 7.
Extend the package	Sections 3–5; then Appendices A–D.
Audit or reproduce a result	Sections 5, 6, and 7; then Appendices D and F.

Conventions used throughout

The economic-geography and urban models share routing, modal choice, congestion, and primitive-cost pass-through, but they have different spatial states and welfare projections. Public transit is therefore a set of modes in the common transport block, not a third spatial model. The guide uses three labels consistently. The *Traditional approach* uses traffic on the improved edge–mode. The *full realized-cost effect* changes the cost users face after congestion. The *Extended approach* is the full primitive-cost effect of the infrastructure-side improvement. Model-ready flows are an input requirement. Vehicle counts, passenger counts, schedules, or tonnage require a documented conversion and balancing step before they enter the package.

Part I Quick start

A new application fails if its policy unit, flow unit, and spatial model are chosen independently. Part I fixes those choices together, documents the source data, constructs model-ready inputs, and produces validated welfare results, maps, and an archived run manifest.

1 Start a new application

To start a new application, the analyst must construct model-ready node and edge–mode flows. The package checks those inputs and computes the derivative; it does not infer economic flows from raw vehicle counts, passenger counts, schedules, or tonnage. The steps below define the policy, prepare the two input tables, validate the baseline, and produce the welfare results.

1.1 Check whether the method fits the policy

The package is designed for marginal improvements to existing, positively used edge–modes. The following screen should be applied before constructing the input files.

Application	Implication
Marginal gross-benefit ranking of existing links with positive baseline use	Directly supported.
Several small improvements to existing links	Use a sparse policy bundle.
New link, new mode, or discontinuous topology change	Outside the interior local derivative.
Large change to an existing link	Use the derivative as a diagnostic and compare it with a nonlinear counterfactual.
Capacity expansion	State how capacity changes the congestion map; a free-flow cost reduction is not automatically equivalent.
Project net present value or benefit–cost ratio	Combine the welfare result with timing, costs, and omitted effects.
Distributional incidence, transfers, or toll revenue	Add the relevant welfare weights and fiscal treatment.

"Local" describes the size of the policy perturbation. It does not describe the adjustment horizon. In the economic-geography model, mobile labor generally gives the derivative a long-run comparative-static interpretation. The urban model likewise permits residence and workplace adjustment.

1.2 Choose the economic setting and policy

Begin by deciding which spatial response the application requires.

Choice	Economic geography	Urban commuting
Movement	Goods between locations	Commuters between residences and workplaces
Node margins	Labor and income	Residents and employment

Choice	Economic geography	Urban commuting
Network flow	Value flow, in the same price basis as income	Commuter flow, normalized by total commuters
Spatial response	Wages and mobile labor	Residence and workplace choices
Typical modes	Road, rail, and water freight	Road, bus, rail or streetcar, and ferry

Public transit is represented by bus, rail, streetcar, or ferry rows in the urban edge–mode file. It does not require a separate transport solver.

Next state the policy. Select a mode, a proportional reduction in primitive generalized cost, and one of three units: a directed edge, a reciprocal physical link, or both. A physical-link experiment requires exactly two opposite policy-mode edges with the same physical-link identifier. Unpaired transit services remain directed policies.

Feature	Status	Use
Economic-geography derivative	Stable reference	U.S. freight and public examples
`ChoiceLogsum` modes	Stable reference	Active multimodal specification
Edge-local congestion	Stable reference	Headline freight specification
Fixed-route decomposition	Verified, memory intensive	Moderate networks after the memory preflight
Endpoint-terminal congestion	Verified extension	Requires an application-specific elasticity
Multimodal urban derivative	Verified extension	Requires commuting inputs that satisfy the urban accounting identities
Seattle workflow	Candidate application	Additional ridership, cost, and calibration evidence required
`ComponentCES`	Legacy only	Historical reproduction

1.3 Install the package and initialize a project

The package API is version 0.2.0 and requires Julia 1.10 or later. Continuous integration tests Julia 1.10 and the current stable release on Linux, macOS, and Windows. Git is needed to clone and update the repository. Running the package requires Julia; building this PDF additionally requires `make`, `latexmk`, Poppler, and `qpdf`. Release-maintainer checks also use Python, but the grid analysis and guide build do not.

Clone the package and instantiate its committed Julia environment:

```
git clone https://github.com/sfuchs-de/TransportNetworkWelfare.jl.git
cd TransportNetworkWelfare.jl
julia --project=. -e 'using Pkg; Pkg.instantiate()'
```

Keep each empirical application outside the package checkout. From the package directory, create either an economic-geography or urban-commuting project:

```
julia --project=. bin/tnw.jl init ../my-freight-network \
  economic_geography
julia --project=. bin/tnw.jl init ../my-commuting-network \
  urban_commuting
```

Omitting the final argument retains the economic-geography default. The command writes the model-specific `config.toml` and seed CSV files, a `sources.toml` provenance ledger, a preprocessing entry point, and a project README:

```
my-network/
```

```
config.toml
sources.toml
data/nodes.csv
data/edge_modes.csv
scripts/prepare_inputs.jl
figures/
```

It refuses to overwrite a nonempty directory. The seed inputs run immediately, providing an installation check before the analyst replaces any data. The preprocessing script marks where a reproducible raw-to-model-ready conversion belongs, but it does not guess how vehicle counts, passenger records, or schedules should be converted into the flows required by the model.

The self-contained examples ordinarily finish in under a minute after Julia precompilation. Runtime depends on hardware and network size. Before `decompose`, `validate` reports the estimated memory for the fixed-route incidence matrix; ordinary `analyze` omits that dense matrix. For the grid used in this guide, that estimate is 0.000373 GiB.

1.4 Replace the template inputs

For economic geography, `nodes.csv` contains

```
node_id,labor,income,longitude,latitude
```

For urban commuting, replace `labor` and `income` with `residents` and `employment`. Coordinates are optional.

Each row of `edge_modes.csv` represents one active mode on one directed edge:

```
edge_id,physical_link_id,origin,destination,mode,flow,
origin_terminal_id,destination_terminal_id,
congestion_elasticity
```

The final three columns are needed only for the corresponding congestion specification. A missing mode row means that the mode is unavailable. Listed flows must be finite and positive. Parallel physical routes require explicit intermediate nodes; they cannot share the same ordered endpoint pair under different edge identifiers.

The baseline must satisfy the accounting identities of the selected model. Economic-geography inputs require

$$\sum_{j,m} \Xi_{ij,m} = \sum_{j,m} \Xi_{ji,m}.$$

Urban inputs instead require

$$l_i^F + \sum_{j,m} \Xi_{ij,m} = l_i^R + \sum_{j,m} \Xi_{ji,m}.$$

Construct any balancing, geographic matching, price conversion, or modal rescaling in a separate preprocessing step. The generic adapter does not perform those transformations silently.

Normalize economic-geography labor and income shares to one. Normalize urban residence and workplace margins separately to one. Divide economic-geography edge value flows by the same world-income aggregate and urban commuter flows by the same total commuter count. Edge traffic itself generally sums to more than one because one shipment or commuter can traverse several edges. Modal shares sum to one within a directed edge.

1.5 Convert standard source data into model inputs

Most applications begin with several data products rather than the two tables read by the package. A node geography supplies the unit of analysis, an origin–destination table supplies residence–workplace or production–consumption margins, a network file supplies available links, and traffic or ridership evidence supplies baseline use. These sources must refer to a common year, geography, population, and flow unit before they can be combined. The table below gives common starting points and the construction each one supports.

Object	Common source	Use in the project
Location boundaries and centroids	National statistical-office boundary files, such as ONS MSOA boundaries (Office for National Statistics, 2021), or a locally maintained GIS layer	Choose node identifiers, calculate centroids, and provide the polygons used for maps. Preserve the source coordinate system until the final export.
Residence–workplace flows	Census or administrative workplace origin–destination tables, such as ODWP01EW for England and Wales (Office for National Statistics, 2023)	Construct urban residence and employment margins. These tables describe OD commuting, not the directed edge flows that the package reads.
Road and rail geometry	Agency GIS files or OpenStreetMap extracts, including the downloadable Greater London files distributed by Geofabrik (Geofabrik, 2026)	Define directed links, reciprocal physical-link identifiers, intermediate route nodes, and map geometry. Record the extract date and license.
Transit service	GTFS, timetable, and station-location feeds, such as TfL's open data (Transport for London, 2026b)	Define stops, links, service availability, travel times, and transfer opportunities. A schedule is not a passenger-flow observation.
Traffic or ridership	Count stations, ticketing records, passenger OD matrices, or freight-value records; for London, TfL network-demand files and ORR station and OD data are useful sources (Transport for London, 2026a; Office of Rail and Road, 2024)	Allocate OD demand across directed edge–modes and check the resulting traffic against independent counts. Station entries and exits alone do not identify segment-level flow.
Generalized cost and congestion	Timetables, speeds, fares, transfer penalties, capacity data, and local engineering or econometric estimates	Construct baseline cost components and choose elasticities. Keep values transferred from another setting in a sensitivity range rather than presenting them as local estimates.

The package automates the calculation after these sources have been converted into one internally consistent baseline. It automates schema checks, route reconstruction, physical-link pairing, the welfare derivative, the decomposition, sensitivity paths, manifests, and standard network figures. It does not automate source-specific geocoding, the valuation of physical freight or passenger counts, OD assignment, or the reconciliation of incompatible populations. Those decisions belong in `scripts/prepare_inputs.jl` and must be described in `sources.toml`. A typical source entry records a stable identifier, URL or local logical path, vintage, license, retrieval date, and SHA-256:

```
[sources.workplace_od]
title = "Census 2021 workplace flows, ODWP01EW"
url = "https://www.nomisweb.co.uk/sources/census_2021_od"
vintage = "2021"
retrieved = "YYYY-MM-DD"
license = "Open Government Licence"
sha256 = "replace-with-downloaded-file-hash"
```

1.6 Prepare geometry and standard figures

Coordinates in `nodes.csv` are enough for a schematic map with straight links. For geographic maps, retain a polygon basemap and a line layer in the same projected coordinate system as the node coordinates. The generic plotter accepts GeoJSON. Each line feature must contain `physical_link_id`; polygon and line features in the basemap need no package-specific attributes. Shapefiles and GeoPackages can be used directly in QGIS, or converted once with GDAL:

```
ogr2ogr -f GeoJSON -t_srs EPSG:27700 \
  data/link_geometry.geojson raw/network_links.gpkg
ogr2ogr -f GeoJSON -t_srs EPSG:27700 \
  data/basemap.geojson raw/msoa_boundaries.gpkg
```

The CRS in these commands is illustrative. The analyst must choose one appropriate to the application and use it consistently across nodes, links, and polygons.

After `analyze` or `decompose`, create a welfare map with:

```
python3 "$PKG/plots/network_example.py" \
  "$APP/data/nodes.csv" "$APP/data/edge_modes.csv" \
  "$APP/output/welfare_physical.csv" \
  "$APP/figures/welfare-map.pdf" \
  --metric primitive_F \
  --link-geometry "$APP/data/link_geometry.geojson" \
  --basemap "$APP/data/basemap.geojson"
python3 "$PKG/plots/hulten_vs_welfare.py" \
  "$APP/output/welfare_physical.csv" \
  "$APP/figures/traditional-versus-extended.pdf"
```

The map command also accepts a PNG output and `-transparent`. Without `-link-geometry`, it draws straight links between node coordinates. The second command draws the Traditional-versus-Extended comparison on common axes and reports Pearson and rank correlations. For publication maps that require labels, insets, or several policy layers, join `welfare_physical.csv` to the line layer by `physical_link_id` in QGIS or another GIS package. The numerical output remains the source of the color scale; the GIS file supplies geometry rather than welfare data.

1.7 Edit the configuration

In the generated `config.toml`, the `spatial_specification` field selects the economic-geography or urban commuting model. The configuration then supplies the corresponding spatial parameters and route curvature.

Set `modal_specification` to `choice_logsum` and supply the modal elasticity η. The congestion block selects edge, endpoint-terminal, composite, or no congestion. The remaining fields define the policy mode, policy unit, shock fraction, output directory, data units, and all transformations applied to the inputs. Do not retain template parameter values merely because they pass validation. Each value needs an empirical source or a stated sensitivity range. Appendix C gives the complete configuration and schema reference.

1.8 Validate, analyze, and decompose

For an application outside the package checkout, use explicit paths:

```
PKG=/path/to/TransportNetworkWelfare.jl
APP=/path/to/my-network
julia --project="$PKG" "$PKG/bin/tnw.jl" validate "$APP/config.toml"
julia --project="$PKG" "$PKG/bin/tnw.jl" analyze "$APP/config.toml"
julia --project="$PKG" "$PKG/bin/tnw.jl" decompose "$APP/config.toml"
```

`validate` checks schemas, identifiers, accounting, route contraction, modal interiority, policy pairing, congestion keys, regularity, and condition numbers. Run `analyze` for the full welfare derivative. Run `decompose` only when the reported fixed-route memory requirement is feasible.

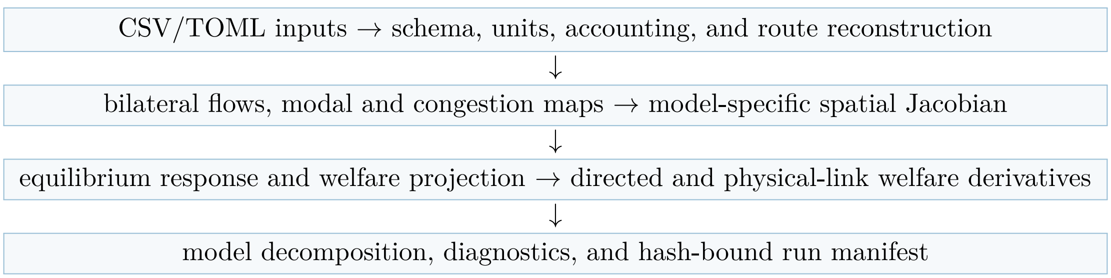

Figure 1: Data-to-result workflow. Application-specific preprocessing precedes the generic package; each package stage checks its output before the next stage uses it.

1.9 Read and preserve the results

`analyze` writes `welfare_directed.csv` and, when requested, `welfare_physical.csv`. `decompose` writes the corresponding `decomposition_*.csv` files. The run also writes `diagnostics.json` and `run_manifest.json`.

A positive reported elasticity is a welfare gain from a primitive-cost reduction. For a small reduction s, the approximate percentage welfare gain is $100s$ times the elasticity. Compare `hulten` with `primitive_F` to see whether equilibrium adjustment changes the traffic-only ranking. Interpret a decomposition only after its identity residuals and condition diagnostics pass.

For example, the highest-ranked grid road link is Row 3, columns 3–4. Its output row reports:

`physical_link_id`	`hulten`	`realized_F`	`primitive_F`
H_r3_c3	0.021040	0.022809	0.010177

Thus the Traditional approach gives the traffic-based elasticity 0.021040, the full realized-cost effect is 0.022809, and the Extended approach gives the primitive-cost effect 0.010177. A one-percent primitive improvement implies $\Delta \log W \simeq 0.00010177$, or approximately 0.010177 percent. Section 5.2 defines the remaining output columns.

Retain the manifest for each result vintage. It records the package commit, configuration, source-ledger, input, and output hashes together with the parameters, transformations, and diagnostics.

1.10 Planning an Elizabeth Line application

Suppose the policy question is which operating Elizabeth Line segment would yield the largest gross welfare gain from a small reduction in generalized travel cost. This is an urban-commuting application because residence and workplace margins differ and commuters can use road and public transport. It is also a marginal application. The package can rank improvements to positive-flow segments of the existing line or evaluate a bundle that improves the operating corridor. It cannot recover the effect of opening the central section from a zero-flow baseline. That question changes the available network and requires a finite counterfactual.

Begin with the urban scaffold:

```
PKG=/path/to/TransportNetworkWelfare.jl
APP=/path/to/elizabeth-line
julia --project="$PKG" "$PKG/bin/tnw.jl" \
  init "$APP" urban_commuting
```

Use 2021 MSOAs or another stable London geography as nodes. ONS ODWP01EW reports counts by residence and workplace area; their row and column sums provide resident and employment margins (Office for National Statistics, 2023). Census Day fell during the March 2021 lockdown, however, so these flows are not a post-opening commuting baseline and must be treated as an initialization or replaced with more representative OD evidence (Office for National Statistics, 2022). The source ledger should state that limitation directly.

TfL timetable and station feeds identify the Elizabeth Line topology, travel times, and stations, while the Greater London OpenStreetMap extract can supply roads and detailed line geometry (Transport for London, 2026b; Geofabrik, 2026). TfL station footfall and ORR station and OD statistics provide post-opening evidence on use (Transport for London, 2026a; Office of Rail and Road, 2024). None of these sources alone is a complete segment-level commuting matrix. The preprocessing script must reconcile their vintages, assign OD commuters to road and transit paths, and report how closely the resulting edge flows reproduce the observed station totals.

The model-ready urban node file contains normalized residents and employment. The edge-mode file contains separate road and rail rows, directed commuter flows normalized by the same total, and terminal identifiers where transfer congestion is modeled. Set `mode = "rail"` in the policy block. Use `unit = "both"` only where the two directions form a valid reciprocal physical link; otherwise retain directed policies. The freight value $\eta = 1.099$ can provide a transparent starting point for sensitivity, but it is not a London estimate. A preferred application would estimate or calibrate modal substitution from London evidence.

Before running the model, edit the following fields in place, retain the other generated input and output entries, and declare the sensitivity values used in the exercise:

```
[input]
mode_order = ["road", "rail"]

[model]
modal_specification = "choice_logsum"
eta = 1.099

[congestion]
specification = "none"

[policy]
mode = "rail"
unit = "both"
shock_fraction = 0.01

[sensitivity]
eta = [0.75, 0.90, 1.099, 1.25, 1.40]
```

The no-congestion specification provides a transparent first run when no London-specific congestion elasticity has been established. An empirical application should add a congestion channel only after documenting its definition, units, and calibration. It should also replace the template's urban spatial parameters rather than treating them as London estimates. When changing the generated composite specification to `none`, delete the generated `[congestion.edge]` and `[congestion.terminal]` subtables.

Once `scripts/prepare_inputs.jl` writes the two CSVs, the full run is:

```
julia "$APP/scripts/prepare_inputs.jl"
julia --project="$PKG" "$PKG/bin/tnw.jl" \
  validate "$APP/config.toml"
julia --project="$PKG" "$PKG/bin/tnw.jl" \
  analyze "$APP/config.toml"
julia --project="$PKG" "$PKG/bin/tnw.jl" \
  decompose "$APP/config.toml"
julia --project="$PKG" "$PKG/bin/tnw.jl" \
  sensitivity "$APP/config.toml"
python3 "$PKG/plots/network_example.py" \
  "$APP/data/nodes.csv" "$APP/data/edge_modes.csv" \
  "$APP/output/welfare_physical.csv" \
  "$APP/figures/elizabeth-line-welfare.pdf" \
  --metric primitive_F \
  --link-geometry "$APP/data/link_geometry.geojson" \
  --basemap "$APP/data/basemap.geojson"
```

The final map ranks marginal improvements to the represented operating segments. A complete appraisal would add construction or operating costs, reliability, crowding, emissions, distributional incidence, and any benefits from new trips outside the maintained commuting baseline.

1.11 Acceptance criteria

The data are ready when their flow unit matches the selected spatial model and congestion elasticities, and when the accounting and route-reconstruction checks pass. The policy is ready when its direction, mode, reciprocal pairing, and shock size describe the intended experiment. The model is ready when each active parameter has a documented source, the calculation remains on a regular branch, and finite-difference or limiting-case tests support any new extension.

The results are ready for interpretation when the sensitivity analysis distinguishes changes in levels from changes in rankings. At that point, archive the manifest with the model-ready inputs so the reported result can be reconstructed from the same data and specification.

To update the package later, record the commit used for the archived results, fast-forward the package checkout, instantiate the environment, rerun `Pkg.test()`, and then revalidate the external project. Part II explains the economics and algorithms behind each of these checks.

Part II Detailed guide: theory and welfare

The quick start states the required inputs and checks without deriving them. Part II supplies that derivation. It defines the two spatial models and their shared transport block, establishes the traffic-based welfare benchmark, and then develops the implicit derivative and decomposition used by the package.

2 Policy experiment and welfare measures

Any welfare comparison is meaningful only relative to a stated policy and baseline. The analysis therefore begins by defining the marginal infrastructure change, then distinguishes the traffic-based and spatial-equilibrium measures and states which benefits they omit.

2.1 The policy experiment

Suppose an analyst can lower the underlying transport cost on one network link. Which link produces the largest economic benefit? A traditional answer multiplies the cost reduction by the traffic using that link. This is the social-savings calculation associated with Fogel (1962) and Fogel (1964). It is simple, data-light, and often informative. It does not require the analyst to model changes in routes, modes, prices, or locations. Such adjustment does not by itself invalidate social savings. The relevant question is whether those responses change exposure to congestion or spatial externalities that give reallocation a first-order welfare value.

The package evaluates a local version of that question. A policy is a small change in the primitive cost of a declared edge–mode. For a directed edge $e = (k, l)$ and mode m, define the primitive log cost as

$$\theta_{klm} = \log \bar{\kappa}_{kl,m}.$$

Here $\bar{\kappa}_{kl,m}$ is the multiplicative transport or generalized commuting cost supplied by infrastructure before endogenous congestion. The policy derivative is

$$-\frac{d \log W}{d\theta_{klm}}.$$

Here W is aggregate welfare in the spatial model defined in Section 3. The reported quantity is the elasticity of that welfare measure with respect to a reduction in primitive edge–mode cost. Multiplying it by 0.01 gives the first-order welfare change from a one-percent improvement. The software reports the elasticity and records the shock size used for this conversion in the run manifest.

The distinction between a directed edge and a physical link matters. A directed policy changes one direction. A physical-link policy changes the two reciprocal directions simultaneously and sums their first-order effects. The package does not treat an unpaired directed edge as bidirectional.

The package evaluates this experiment in either the economic-geography or urban-commuting model. Section 3 defines their distinct economic adjustments and the transport equations they share. The welfare measures below apply to both models, with traffic expressed in the units appropriate to each.

2.2 Traditional and extended measures

Let $\Xi_{kl,m}$ be traffic on the edge–mode in the normalization of the selected spatial model. It is a world-income share in the economic-geography model and a share of the commuter flow in the urban model. The Traditional approach is

$$H_{klm} = \Xi_{kl,m}.$$

It is a welfare elasticity, not a raw vehicle count. In the economic-geography model, multiplying by world income converts the traffic share to currency units before applying a small proportional cost change. In the urban model, the conversion depends on the generalized commuting-cost units used to construct the baseline.

Two full-model effects add the responses omitted from this traffic measure. The full realized-cost effect changes the cost users face after congestion. The Extended approach reports the full primitive-cost effect of the infrastructure-side improvement. Pass-through separates these experiments because induced traffic can attenuate the user-facing cost reduction. Both effects value the interaction of route, modal, congestion, and spatial responses with pre-existing wedges. In the urban model, the spatial responses are residence and workplace choices rather than trade, wages, and mobile labor.

The economic-geography formula developed in Section 4.2 can be written as four factors:

$$\text{Economic-geography Extended benefit} = \text{traffic} \times \text{primitive-to-user pass-through} \times \text{common spatial scale} \times \text{endpoint access value.}$$

Traffic gives the direct scale of the affected activity. Pass-through measures the share of the primitive improvement that reaches users after transport feedback. The spatial scale depends on the common externality parameters. The endpoint multipliers measure the marginal welfare value of improving access at the link's origin and destination after the rest of the spatial economy adjusts.

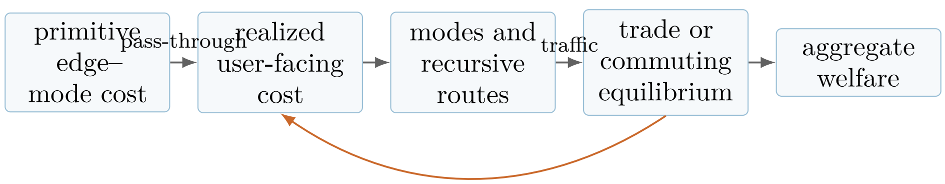

Figure 2: Policy transmission and welfare. The policy changes a primitive infrastructure cost. Congestion and modal adjustment determine the realized cost users face; routes transmit that change to trade or commuting flows; the selected spatial model maps the resulting allocation into welfare.

Local analysis separates the model's economic response from assumptions about project scale and permits a common ranking of many candidate links. The adjoint method also computes all link derivatives from one equilibrium linearization. These advantages come with a clear limit: a large project can leave the neighborhood in which the derivative is accurate. Section 5.3 explains how to compare the first-order prediction with finite changes.

2.3 Benefits included in the welfare measure

The welfare measure covers only effects represented in the selected model: resource or generalized costs of movement, route and mode substitution, congestion, spatial externalities, and changes in economic location. Construction, maintenance, financing, safety, environmental, neighborhood, displacement, segregation, and distributional effects remain separate entries in a complete appraisal.

Practical caution. Do not label the output "net social benefit" unless all material project costs and externalities have been brought into the welfare criterion. In the paper application, "model-implied economic benefit" is the accurate description.

The package therefore complements conventional transport appraisal rather than replacing it. Conventional analysis supplies engineering costs, demand evidence, safety effects, and local externalities. The spatial statistic shows the contribution of observed network use and equilibrium adjustment to the benefit side of that ledger. For conventional appraisal, see Small et al. (2024) and de Palma et al. (2011); for spatial effects, see Redding and Turner (2015) and Donaldson (2025).

2.4 Using the elasticity in project appraisal

Let $E_p = -d \log W / d\theta_p$ be the reported benefit elasticity and let $s_p > 0$ be a small proportional primitive-cost reduction. The local welfare change is

$$\Delta \log W \simeq E_p s_p.$$

This statement is invariant to the currency used in the input files. A money-metric conversion requires an additional normalization. If W is measured in units for which the chosen income aggregate Y_0 is the relevant baseline expenditure function, a first-order equivalent-variation approximation is

$$B_0 \simeq E_p s_p Y_0.$$

For economic geography, the paper normalizes value traffic by world income, so the corresponding baseline aggregate is $Y_0 = Y^W$, measured over the same period as the income and traffic data. For an urban application, the commuter-share normalization alone does not supply a currency conversion; the analyst must specify the generalized-cost units and the money-metric welfare normalization.

The package reports E_p, not an annualized project value. Treat B_0 as an annual flow only when the baseline income, traffic, and welfare normalization all refer to that annual period. A conventional appraisal then adds omitted benefits and subtracts construction, maintenance, financing, and other costs at their own dates:

$$\text{NPV} = \sum_t \frac{B_t + B_t^{\text{other}} - C_t}{(1+r)^t}.$$

Transfers such as toll revenue require separate treatment from resource costs. Large projects also require a nonlinear counterfactual or other evidence that the local approximation remains accurate at the proposed scale.

The quick start in Section 1 implements these requirements as a data-to-results workflow. The detailed argument now turns to the two spatial models and the transport block they share. Section 4 then establishes when traffic alone is exact and derives the correction when wedges matter, before Section 5 translates the formulas and required inputs into data and code.

3 Spatial models and shared transport technology

The policy derivative requires two components. A transport block maps edge–mode costs into routes, modal traffic, and congestion. A spatial model maps the resulting bilateral costs into economic activity and welfare. The package shares the transport block across applications and supplies separate economic-geography and urban equilibrium systems.

3.1 Two spatial models, one transport block

The two models start from the same directed network and the same primitive edge–mode policy. They use the negative-power modal logsum, recursive route aggregation, flexible and fixed route or mode comparisons, edge and terminal congestion, and primitive-cost pass-through. They differ in the economic flows assigned to the network and in the equilibrium response projected onto welfare.

Model element	Economic geography	Urban commuting
Economic movement	Goods shipped between locations	Commuters traveling between residences and workplaces
Location margins	Labor and income; source and destination income margins coincide	Residence l^R and workplace l^F margins can differ

Model element	Economic geography	Urban commuting
Network flow	Value traffic normalized by world income	Commuter traffic normalized by total commuters
Spatial response	Wages and the location of mobile labor	Residence and workplace choices
Welfare projection	Common real income or utility under mobility	$-d\log\varpi_U/\theta_U$
Route curvature	$\sigma-1$	Commuting heterogeneity θ_U
Typical modes	Road, rail, and water freight	Road, bus, rail or streetcar, and ferry

The table separates economic interpretation from transport technology. Road, bus, rail, streetcar, and ferry enter as modes in the urban transport network; public transit is not a third spatial model. Adding a mode changes the costs and traffic available to commuters without changing the residence–workplace equilibrium. Conversely, using the same road network in the two models does not make their welfare derivatives identical because their spatial states and welfare rows differ.

3.2 Economic-geography model

The economic-geography closure turns bilateral transport costs into trade, production, consumption, and mobile labor. We first define those location-level objects, then derive bilateral trade and the transformed equilibrium used by the welfare derivative.

3.2.1 Production, consumption, and mobility

The economic-geography model in Allen et al. (2026) has N locations. Location i has labor L_i, wage w_i, productivity A_i, and amenity u_i. A differentiated variety is produced with labor under constant returns and perfect competition. Consumers have CES preferences with elasticity $\sigma > 1$. With P_i denoting the CES price index, indirect utility in location i is proportional to

$$U_i = \frac{w_i u_i}{P_i}.$$

Local activity can affect productivity and amenities:

$$A_i = \bar{A}_i L_i^{\alpha}, \qquad u_i = \bar{u}_i L_i^{\beta}.$$

The parameter α is the elasticity of productivity with respect to local labor, while β is the corresponding amenity elasticity. Labor is perfectly mobile across locations, aggregate labor is fixed, and each worker supplies one unit inelastically. Spatial equilibrium equalizes utility across occupied locations and clears goods and labor markets.

These assumptions determine the economic-geography welfare formula. Perfect mobility equalizes utility, while fixed aggregate labor and labor-only production exclude further adjustment through factor income or labor supply. The urban model below replaces these assumptions with separate residence and workplace choices.

3.2.2 Bilateral trade and market access

Let τ_{ij} be the composite iceberg cost of delivering a good from i to j. CES demand and the production assumptions imply

$$X_{ij} = W^{1-\sigma}\left(\frac{A_i u_j}{\tau_{ij}}\right)^{\sigma-1} w_i^{1-\sigma} w_j^{\sigma} L_j.$$

Here W is aggregate welfare in the normalization used by the paper. The network determines τ_{ij}; the economic-geography equilibrium determines wages, labor, productivity, and amenities.

Direct computation of the bilateral cost matrix would obscure the network structure. Following Allen et al. (2026), let Π_i denote the outward price-index analogue and define the inward and outward market-access stocks as

$$\mathcal{P}_i = P_i^{1-\sigma}, \qquad \mathcal{Q}_i = \Pi_i^{1-\sigma}.$$

Define the local terms

$$a_i^Q = W^{1-\sigma} u_i^{\sigma-1} w_i^{\sigma} L_i, \qquad a_i^P = A_i^{\sigma-1} w_i^{1-\sigma}.$$

The two market-access stocks satisfy oppositely oriented local recursions:

$$\mathcal{Q}_i = a_i^Q + \sum_{j \in \mathcal{N}^+(i)} \kappa_{ij}^{1-\sigma} \mathcal{Q}_j, \tag{3.1}$$

$$\mathcal{P}_i = a_i^P + \sum_{j \in \mathcal{N}^-(i)} \kappa_{ji}^{1-\sigma} \mathcal{P}_j. \tag{3.2}$$

Thus current-node access is divided between local absorption and continuation through adjacent edges. These sparse recursions replace the dense bilateral cost matrix. Allen et al. (2026, Appendix A) derives them from goods market clearing; Appendix A of this guide shows how the terms in these recursions enter the code.

3.2.3 Trade exposure and the transformed equilibrium

Define economic-geography trade exposure by

$$\mathcal{T}_i^G = \frac{\mathcal{P}_i \mathcal{Q}_i}{Y^W}, \qquad Y^W = \sum_i w_i L_i.$$

The access recursions divide this stock into local absorption and continuation. Define

$$s_i^x = \frac{a_i^Q}{\mathcal{Q}_i}, \qquad \mu_{ij} = \frac{\kappa_{ij}^{1-\sigma} \mathcal{Q}_j}{\mathcal{Q}_i}, \qquad s_i^y = \frac{a_i^P}{\mathcal{P}_i}, \qquad \lambda_{ij}^{in} = \frac{\kappa_{ij}^{1-\sigma} \mathcal{P}_i}{\mathcal{P}_j}.$$

Here μ_{ij} is the outward continuation share on $i \to j$, while λ_{ij}^{in} is the same edge's share of destination-j exposure. Directed edge traffic factors as

$$\Xi_{ij} = \mu_{ij} \mathcal{T}_i^G = \lambda_{ij}^{in} \mathcal{T}_j^G = \frac{\kappa_{ij}^{1-\sigma} \mathcal{P}_i \mathcal{Q}_j}{Y^W}.$$

Outgoing and incoming traffic totals therefore satisfy

$$O_i = (1 - s_i^x) \mathcal{T}_i^G, \qquad I_i = (1 - s_i^y) \mathcal{T}_i^G,$$

for local absorption shares s_i^x and s_i^y. The package checks that outgoing and incoming recovery of $\mathcal{T}_i^G$ agree.

Allen et al. (2026, Appendix A) work with the transformed variables

$$x_i = w_i^{\sigma} L_i^{1+\beta(\sigma-1)}, \qquad y_i = W^{\sigma-1} w_i^{1-\sigma} L_i^{\alpha(\sigma-1)}.$$

This transformation separates sparse market-access propagation from the aggregate labor constraint. Define

$$e = 1 + \beta(\sigma - 1) + \alpha\sigma, \qquad \rho = \frac{1 + \alpha + \beta}{e}.$$

The inverse map from (x, y) to wages, labor, and welfare requires $e \neq 0$; the welfare normalization also requires $1 + \alpha + \beta \neq 0$. Stack the transformed log changes in the $2N$-element column vector

$$z_G = (d \log x_1, \ldots, d \log x_N, d \log y_1, \ldots, d \log y_N)^\top.$$

The route, mode, and congestion maps below translate a primitive edge–mode shock into the aggregate edge-cost changes entering the residual system for z_G.

Implementation. The economic-geography coefficient map is implemented by `AdjointRSUE.coefs`. `AdjointRSUE.assemble_J` constructs the no-congestion spatial Jacobian. The generic builder then adds the common route, modal, and congestion feedback.

3.3 Urban commuting: residences and workplaces

The urban specification follows the residence–workplace model in Allen and Arkolakis (2022). Let l_i^R and l_i^F denote normalized residence and workplace masses. A location can contain many residents but few jobs, or many jobs but few residents. Commuting flows therefore have distinct origin and destination margins.

The urban accounting stock satisfies

$$\mathcal{T}_i^U = l_i^F + \sum_j \Xi_{ij} = l_i^R + \sum_j \Xi_{ji}.$$

The first equality combines workplace absorption with outgoing route traffic; the second combines residence absorption with incoming traffic. It is the urban counterpart to trade exposure, but it is measured in normalized commuter flows rather than shares of world income. The loader rejects a baseline when the two expressions disagree beyond the declared tolerance.

The differentiated urban state and welfare change are

$$z_U = (d \log l^R, d \log l^F, d \log \varpi_U)^\top, \qquad d \log W = -\frac{d \log \varpi_U}{\theta_U}.$$

We write the urban equilibrium normalizer as ϖ_U to distinguish it from the primitive-cost pass-through χ_{klm}. The implementation retains `chi_U` as the historical name of this final state component. The parameter $\theta_U > 0$ governs commuting heterogeneity and supplies the route curvature. The urban parameters α and β enter the workplace- and residence-side response coefficients. The resulting Jacobian J_U^0, aggregate-cost loading $B_{\kappa,U}$, and welfare row

$$q_U^\top = (0, \ldots, 0, -1/\theta_U)$$

are specific to the residence–workplace equilibrium. They cannot be obtained by relabeling the trade equations.

Public transit enters through the shared edge–mode network. A directed edge can carry road, bus, rail, streetcar, or ferry traffic, and the modal logsum determines how their generalized costs combine. The urban model then maps the induced commuting-cost changes into residence, workplace, and welfare responses. Subsection 7.2 documents the synthetic urban test and the Seattle data candidate.

3.4 Interface with the transport block

For spatial model $c \in \{G, U\}$, write its no-congestion residual as

$$G_c^0(z_c, \kappa) = 0.$$

Each model supplies three inputs to the shared computation: a spatial Jacobian $J_{c,0}$, a loading $B_{\kappa,c}$ from aggregate edge costs into the residuals, and a welfare row q_c. The next subsection constructs the route, modal, and congestion derivatives combined with these inputs. Subsection 4.2 then derives welfare after eliminating the transport block.

3.5 Routes, modes, and congestion

The spatial models defined above take bilateral generalized costs as inputs, but the data describe directed edges and modes. The transport block connects these levels. It aggregates edges into routes, reconstructs baseline traffic, combines modal costs, and maps traffic changes into congestion.

3.5.1 Recursive path aggregation

Each spatial model takes bilateral transport costs as inputs. Those costs must therefore be related to the directed edges and modes observed in the network before either model can compute a welfare response. A path from i to j is a sequence of directed edges. If κ_{kl} is an iceberg cost factor, a path cost is the product of its edge costs. Equivalently, generalized costs add after taking logs. With a hard minimum, the bilateral cost is

$$\widetilde{\tau}_{ij} = \min_{p \in \mathcal{P}^{ij}} \prod_{(k,l) \in p} \kappa_{kl}.$$

The model instead smooths this choice over feasible routes (Allen et al., 2026):

$$\tau_{ij}^{-\omega} = \mathbf{1}\{i = j\} + \sum_{k \in \mathcal{N}^{+}(i)} \kappa_{ik}^{-\omega} \tau_{kj}^{-\omega}.$$

The symbol ω denotes route curvature. For the economic-geography theorem, $\omega = \sigma - 1$; for the urban model, $\omega = \theta_U$, its commuting heterogeneity parameter. The package rejects an independent economic-geography route curvature because the paper's derivative is proved for the common-curvature case.

The recursion can be written using a substochastic continuation matrix K. For a spectral radius below one, the route resolvent is

$$T = (I - K)^{-1} = I + K + K^2 + \cdots .$$

The series shows why the inverse has a route interpretation: it sums weighted walks of arbitrary length until traffic is absorbed. Differentiating the same resolvent measures how an edge-cost change alters all origin–destination pairs whose weighted walks use that edge.

3.5.2 Reconstructing baseline routes

The package starts from directed edge traffic and the location margins appropriate to the selected model, not an unobserved path table. It reconstructs the route kernel from continuation shares and obtains bilateral flows as

$$X^{OD} = \operatorname{diag}(o)\,(I - K)^{-1}\,\operatorname{diag}(a).$$

Here o is the source margin and a the destination absorption vector in the relevant model. For economic geography the source and destination margins coincide with normalized income. For urban commuting they are constructed from the distinct residence and workplace margins described in Section 3.

The reconstruction is accepted only when the route kernel is contractive and the absorption probabilities add to one. The resulting bilateral flows must match the declared row and column margins, while the traffic reconstructed on each directed edge must match the input. In addition, each edge affected by a policy must carry positive baseline flow in the relevant mode.

The fixed-route case stores baseline origin-destination edge incidence. The flexible-route case differentiates the resolvent. Both are evaluated at the same observed baseline, which makes their difference an exact local route wedge rather than a comparison of two separately calibrated models.

Route reconstruction treats κ_e as an aggregate edge cost. The modal block determines how the costs and traffic of the modes available on that edge combine into κ_e.

3.5.3 Choice-consistent modal aggregation

Each directed edge can carry several modes. The active paper specification adapts the multimodal structure in Fuchs and Wong (2026) and uses

$$\kappa_e = \left(\sum_m \kappa_{e,m}^{-\eta}\right)^{-1/\eta}, \qquad s_{e,m} = \frac{\kappa_{e,m}^{-\eta}}{\sum_n \kappa_{e,n}^{-\eta}}.$$

The share is both the cost elasticity $\partial \log \kappa_e / \partial \log \kappa_{e,m}$ and the modal traffic share in this logsum. A lower modal cost raises its share. A larger η makes modal allocation more responsive to relative costs. Baseline modal shares are model-ready data inputs, with mode-specific shifters rationalizing those shares. The parameter η governs the local response to a cost change; it does not overwrite the baseline allocation.

The active economic-geography and urban specifications use `ChoiceLogsum`. `ComponentCES` is **legacy-only**; Appendices A and F document that earlier positive-power implementation.

Modal substitution changes edge-mode traffic. When realized costs depend on that traffic, the modal and congestion responses must be solved jointly.

3.5.4 Congestion maps

Realized edge-mode cost can differ from the primitive infrastructure cost:

$$\log \kappa_{e,m} = \log \bar{\kappa}_{e,m} + \gamma_{e,m} \log \Xi_{e,m}$$

for edge-local congestion. The headline application sets $\gamma_{e,\text{road}} = 0.092$, following Allen and Arkolakis (2020), and sets other headline congestion elasticities to zero.

The package also allows endpoint-terminal congestion. In that extension, an edge-mode cost depends on total affected traffic at the declared origin and destination terminals. This specification applies to ports, rail terminals, or transit hubs, but it is not part of the paper's main theorem application.

Section 4.2 differentiates the route, modal, and congestion equations jointly and eliminates their responses before solving the spatial system. The code checks that transport system before adding its feedback to the spatial Jacobian; a singular system stops the calculation.

> **Practical caution.** Traffic must be measured in the units implied by the congestion elasticity. An economic-geography value-flow elasticity cannot be applied automatically to a vehicle count. An urban commuter-flow elasticity cannot be applied to scheduled service frequency without a documented mapping. A BPR elasticity also requires a declared baseline volume-to-capacity ratio.

These transport equations determine how shippers or commuters reallocate after a cost change. At an efficient baseline, the envelope theorem implies that even substantial reallocation need not affect welfare at first order. The next section establishes when traffic alone measures welfare exactly and then derives the correction produced by congestion or spatial externalities.

4 Traffic and spatial-equilibrium welfare

The preceding section maps a primitive link cost into bilateral costs and edge–mode traffic. We can now ask when that traffic is itself the welfare statistic and when equilibrium adjustment adds a first-order correction. The answer begins with the efficient social-savings benchmark.

4.1 Traditional approach: social savings

The Traditional approach first treats observed traffic as an accounting measure. The envelope argument that follows explains when the same expression also equals the full marginal welfare effect.

4.1.1 The accounting statistic

Social savings holds observed edge–mode use fixed and asks how much less it would cost to move the same goods or commuters after an infrastructure improvement. For a marginal proportional reduction in one edge–mode cost, homothetic preferences give the welfare elasticity

$$H_{klm} = \Xi_{kl,m}.$$

The expression is attractive because it needs only the baseline value flow. It does not require solving a counterfactual spatial equilibrium.

Holding the shipment pattern fixed describes how the statistic is computed, not the equilibrium content of the result. In the efficient case stated below, it is the exact first-order general-equilibrium effect even though trade, routes, modes, prices, and locations adjust.

Proposition 4.1 (Efficient social savings in economic geography)**.** *Suppose $\alpha = \beta = 0$ and all congestion elasticities are zero. At an interior efficient equilibrium,*

$$-\frac{d \log W}{d \log \bar{\kappa}_{kl,m}} = \Xi_{kl,m}.$$

This is the formal result proved for the paper's economic-geography model. The urban model has a different residual system and welfare row. Its one-mode efficient limit must nevertheless reproduce the corresponding edge-mode commuter-traffic statistic; the package treats that equality as a separate urban regression test rather than reusing the trade-model proof.

4.1.2 Envelope logic

The result is an application of the envelope theorem and the network version of Hulten's theorem (Hulten, 1978; Baqaee and Farhi, 2019). At an efficient allocation, the planner is locally indifferent to reallocations that preserve feasibility. A lower transport cost directly saves resources. The induced changes in routes, modes, trade or commuting, and location are responses along the feasible set. Their first-order welfare effects vanish at the optimum.

The statement is stronger than saying that adjustment is quantitatively small. Adjustment can be large. What vanishes is its first-order contribution to welfare at an efficient baseline. This distinction explains why adding a rich spatial equilibrium to the model does not, by itself, invalidate social savings. Only wedges that give the induced reallocations a first-order welfare value require a correction.

4.1.3 Failures of the efficient result

Congestion and spatial externalities break the envelope result because a transport improvement changes exposure to pre-existing wedges. Route diversion can relieve or worsen congestion. Relocation can raise productivity through agglomeration or affect amenities through dispersion. These responses need not vanish at first order.

The result also depends on the assumptions used to define welfare and mobility. The paper formally proves it under perfect labor mobility and the stated aggregate welfare criterion. Similar envelope logic can apply under other mobility settings only when the allocation is efficient in the enlarged problem and welfare is aggregated consistently. Fixed factors, taxes, heterogeneous households, or endogenous labor supply can add incidence terms.

4.1.4 Uses of the traffic measure

The package retains the Hulten statistic even when the full model is distorted. It provides the traffic-based Traditional approach and an exact unit test in the efficient limit. It also supplies the reference point for decomposing why the full primitive-cost result differs from traffic alone.

If the efficient fixture does not collapse to traffic, the calculation has a sign, normalization, or indexing error. If the distorted model remains close to Hulten, the data and calibration produced that result; the model did not impose it.

> **Implementation.** The output column `hulten` is the policy-mode traffic share in the selected model's normalization: world-income-normalized value traffic for economic geography and normalized commuter traffic for the urban model. The full realized- and primitive-cost elasticities are reported separately. Physical-link results sum opposite directions before computing link-level comparisons.

Once congestion or spatial externalities break the envelope result, welfare depends on the equilibrium response rather than traffic alone. The next subsection obtains that response by differentiating the joint equilibrium system.

4.2 Extended approach: implicit derivative and adjoint

The efficient case removes the need to trace equilibrium adjustment. Externalities and congestion remove that simplification because reallocations then have first-order welfare effects. The Extended approach differentiates the equilibrium system and projects the resulting state response onto welfare.

4.2.1 Need for an implicit derivative

With externalities or congestion, the welfare effect of a link improvement is not exhausted by its direct resource saving. The improvement changes realized costs, routes, and modal allocation. It also changes either trade, wages, and mobile labor or commuting, residence, and workplace choices. These outcomes are determined jointly. Differentiating one equation at a time while holding the remaining equilibrium quantities fixed would omit the feedback that motivates the exercise.

Let $c \in \{G, U\}$ index the economic-geography or urban model. Let $z_c \in \mathbb{R}^{n_c}$ be the column vector of transformed equilibrium variables defined in Section 3, and write its n_c residual equations as

$$G_c(z_c, \theta) = 0. \tag{4.1}$$

For economic geography, z_G contains changes in the transformed x and y variables. For the urban model, z_U contains changes in residence, workplace, and ϖ_U. The column vector θ contains primitive edge–mode log costs. Around an interior differentiable baseline,

$$J_c \, dz_c + G_{c,\theta} \, d\theta = 0, \qquad J_c = G_{c,z}.$$

If J_c is nonsingular, the Implicit Function Theorem gives

$$\frac{dz_c}{d\theta} = -J_c^{-1} G_{c,\theta}. \tag{4.2}$$

Let $q_c^\top = d \log W / dz_c$ be the model-specific welfare row and let $b_{c,p} = G_{c,\theta_p}$ be the residual response to policy p. Substituting (4.2) into the welfare differential gives

$$-\frac{d \log W}{d\theta_p} = q_c^\top J_c^{-1} b_{c,p}. \tag{4.3}$$

The sign convention makes a reduction in cost a positive benefit when the reported elasticity is positive.

4.2.2 The adjoint solve

There may be hundreds or thousands of candidate links but only one welfare criterion for the selected model. Solving $J_c x_p = b_{c,p}$ separately for each policy repeats most of the linear algebra. The adjoint method instead solves

$$J_c^\top \psi_c = q_c$$

once and evaluates

$$E_{c,p} = \psi_c^\top b_{c,p}. \tag{4.4}$$

Since $\psi_c^\top = q_c^\top J_c^{-1}$, this expression is identical to (4.3). After the adjoint solve, each additional policy requires only sparse matrix products.

The package does not form an explicit matrix inverse. It factorizes the relevant systems and uses linear solves. This matters for numerical stability and for moderately large sparse networks.

4.2.3 Spatial and transport blocks

Each model has a sparse spatial block and uses the common route–mode–congestion block. For transport case S, stack active edge–mode log-traffic changes in ξ and realized edge–mode log-cost changes in h. The differentiated transport block is

$$d\xi = Q_{z_c}^S dz_c + C^S dh, \qquad dh = d\theta + \Gamma^S d\xi.$$

The matrix $Q_{z_c}^S$ maps spatial changes into traffic, C^S maps realized costs into traffic through routes and modes, and Γ^S maps traffic into congestion costs. Substitution of the traffic equation into the cost equation gives

$$dh = (I - \Gamma^S C^S)^{-1} \left(d\theta + \Gamma^S Q_{z_c}^S dz_c \right).$$

This expression separates the two roles of the transport block. The term in dz_c feeds congestion back into the equilibrium Jacobian. The term in $d\theta$ transmits the primitive policy through modal, route, and congestion responses and supplies the primitive forcing vector. A realized-cost experiment instead treats the change in h as the user-facing policy input.

Let $J_{c,0}$ be model c's no-congestion spatial Jacobian. $B_{\kappa,c}$ loads aggregate directed-edge costs into its residuals, while S_{agg} aggregates edge–mode costs. Setting $d\theta = 0$ in the preceding expression isolates state feedback. Substituting that term into the spatial residuals gives

$$J_{c,S} = J_{c,0} + B_{\kappa,c} S_{\text{agg}} (I - \Gamma^S C^S)^{-1} \Gamma^S Q_{z_c}^S. \tag{4.5}$$

Equation (4.5) is the Schur complement obtained by eliminating the differentiated transport quantities. The transport matrices have the same definitions in both models; their entries depend on each model's margins and state responses. Replacing J_c in (4.3) with $J_{c,S}$, and using the forcing generated by the $d\theta$ term above, gives the welfare derivative for transport case S. Appendix B gives this computational sequence, and the test suite compares the reduced and uneliminated systems.

4.2.4 Realized- and primitive-cost effects

The output contains two full-model derivatives. A realized-cost shock changes the cost users face after congestion:

$$\vartheta_{klm} = \log \kappa_{kl,m}.$$

A primitive-cost shock changes the infrastructure component:

$$\theta_{klm} = \log \bar{\kappa}_{kl,m}.$$

With congestion, these are different experiments. Traffic induced by a primitive improvement can attenuate the fall in the cost users face. The CSV field `realized_F` reports the first effect; `primitive_F` is the Extended approach and the headline policy measure.

To separate modal exposure from the transport system's response, define

$$\zeta_{klm} = \left. \frac{\partial \log \kappa_{kl}}{\partial \theta_{klm}} \right|_{dz_c = 0}, \qquad \chi_{klm} = \frac{\zeta_{klm}}{s_{kl,m}}. \tag{4.6}$$

The derivative is taken through the route–mode–congestion block while holding the selected spatial state fixed. With no congestion feedback, $\zeta_{klm} = s_{kl,m}$ and $\chi_{klm} = 1$; deviations from one measure attenuation or amplification within the transport block. The spatial equilibrium then responds through J_c^{-1}. The implementation constructs the full primitive forcing rather than assigning χ separately.

For a nonzero realized welfare derivative, the output reports the welfare-effective ratio

$$\chi_e^{eff} = \frac{E^{\theta}_{e,F}}{E^{r}_{e,F}}.$$

If the denominator is numerically zero, the ratio is missing and the pass-through contribution is still reported in levels. A zero derivative is a valid result, not an exception.

Measure	Comparison	Interpretation
χ_{klm}	$\zeta_{klm}/s_{kl,m}$	Elasticity of the aggregate edge cost with respect to a primitive edge–mode cost, holding the spatial state fixed.
χ_e^{eff}	$E^{\theta}_{e,F}/E^{r}_{e,F}$	Ratio of two full-equilibrium welfare derivatives. It can include nonlocal congestion responses and is undefined when the realized derivative is zero.

4.2.5 Economic-geography factorization

The adjoint formula computes the declared policy derivatives efficiently, but it does not by itself show why links with similar traffic can have different welfare effects. For the paper's economic-geography specification with common route curvature and edge-local congestion, the expression in (4.3) can be factored into traffic, primitive-cost pass-through, a common externality scale, and link-specific market-access multipliers. Using the regularity scalar e from Section 3, Allen et al. (2026) define

$$\rho = \frac{1+\alpha+\beta}{1+\beta(\sigma-1)+\alpha\sigma} = \frac{1+\alpha+\beta}{e}.$$

Let $\mathcal{M}_i^{in}$ and $\mathcal{M}_i^{out}$ be the adjoint market-access multipliers normalized by trade exposure $\mathcal{T}_i^G$. The primitive edge–mode welfare elasticity is

$$-\frac{d\log W}{d\theta_{klm}} = \chi_{klm}\rho\Xi_{kl,m}\left(\mathcal{M}_k^{in} + \mathcal{M}_l^{out}\right). \tag{4.7}$$

The factorization is algebraically equivalent to the economic-geography adjoint derivative. Each term answers a different part of the welfare calculation. Traffic $\Xi_{kl,m}$ measures the economic activity directly exposed to the improvement. The externality scale ρ applies the same transformation of welfare and spatial responses uniformly across links for a given parameter vector. Pass-through χ_{klm} measures how much of the primitive change reaches the aggregate edge cost after modal and congestion responses. The inward and outward market-access multipliers then determine how the shock propagates through equilibrium at the two endpoints.

Appendix A maps these factors to their code counterparts. The paper's Appendix A supplies the full residual-system and adjoint derivation; Appendix B here records the corresponding order of computation.

The factorization reduces to the Traditional approach in the efficient case. At an efficient interior baseline with no congestion,

$$\alpha = \beta = 0, \qquad \Gamma = 0 \implies \rho = 1, \quad \chi_{klm} = 1, \quad \mathcal{M}_k^{in} + \mathcal{M}_l^{out} = 1.$$

Substitution into (4.7) gives

$$-\frac{d\log W}{d\theta_{klm}} = \Xi_{kl,m},$$

the result in Proposition 4.1. Congestion or spatial externalities change at least one of the three nontraffic factors and create the correction measured by the Extended approach.

4.2.6 Urban welfare projection

The economic-geography factorization relies on its trade equilibrium and therefore should not be relabeled as an urban formula. The urban model uses the same adjoint calculation and primitive transport forcing, but it does not reuse ρ or the trade-model market-access multipliers. Its exact derivative is

$$E_{U,p} = q_U^\top J_{U,S}^{-1} b_{U,p}^\theta, \qquad q_U^\top = (0,\ldots,0,-1/\theta_U).$$

The residence–workplace Jacobian determines how a commuting-cost shock changes l^R, l^F, and ϖ_U. Flexible routes, modal substitution, edge or terminal congestion, and primitive-cost pass-through enter $J_{U,S}$ and $b_{U,p}^\theta$ through the same transport operators as in economic geography. The urban result is therefore methodologically parallel to Proposition 2 without being its scalar trade-exposure factorization.

4.2.7 Status of inputs and computed quantities

The derivative combines model-ready inputs with reconstructed, calibrated, and computed quantities. Keeping those categories separate prevents an accounting identity from being described as external validation.

Status	Inputs or quantities	Empirical responsibility
Model-ready inputs	Location margins, edge–mode flows, edge and terminal identifiers, and the policy-to-edge definition	Observe directly or construct through a documented preprocessing step in the units required by the selected model.
Reconstructed	Absorption and continuation shares, route kernel, bilateral flows, and exposure stocks	Verify accounting, contraction, margins, and recovery of directed-edge traffic.
Calibrated	Spatial, route, modal, and congestion elasticities	Cite an estimate, report a transfer from another setting, or provide a sensitivity range.
Computed	Jacobians, forcing matrices, adjoint vectors, welfare elasticities, and decomposition terms	Check solve residuals, conditioning, finite differences, limiting cases, and decomposition identities.

The exposure stocks show why the distinction matters. In economic geography, outgoing or incoming traffic gives

$$\mathcal{T}_i^G = \frac{O_i}{1-s_i^x} = \frac{I_i}{1-s_i^y}.$$

For the urban model, commuting exposure satisfies

$$\mathcal{T}_i^U = l_i^F + O_i = l_i^R + I_i.$$

Both are model identities. Empirical implementation requires flows that satisfy the selected model's accounting identities. The package reports exposure disagreement rather than hiding it through an automatic balancing step.

Implementation. The package constructs the common transport response, then combines it with either the economic-geography or residence–workplace equilibrium. The infrastructure shock vector is passed through modes, routes, and congestion before the model-specific welfare formula is applied. Appendices A and B give the corresponding function names and matrix operations.

4.2.8 Regularity and failure

The transport derivative requires an interior baseline with positive flows on active edge–mode pairs, a contractive route kernel, and positive exposure stocks. The economic-geography transformation further requires $e \neq 0$ and $1 + \alpha + \beta \neq 0$. The urban transformation instead requires $\theta_U > 0$ and a nonsingular mapping between residence and workplace allocations.

For either spatial model, the modal-congestion and spatial systems must be nonsingular, and their condition numbers must remain below the configured gate. The software fails explicitly when these conditions are violated. It does not use a ridge term to turn an unidentified derivative into a number.

The derivative measures the welfare effect of one policy. Comparing cases built from the same baseline and policy then shows which model mechanisms account for the difference from the Traditional approach.

4.3 Decomposing the welfare difference

The adjoint calculation gives the full derivative but does not separate its departure from traffic into distinct economic sources. To separate those sources, the package compares models that share the observed baseline and policy forcing but selectively hold congestion, modes, or routes fixed.

4.3.1 Common-baseline comparisons

The full primitive-cost welfare elasticity can differ from traffic for several reasons. To isolate them, the package evaluates each alternative at the same observed baseline and under the same policy forcing, following the inverse-Jacobian logic in Fuchs and Wong (2026). Recalibrating each alternative would mix the effect under study with a change in starting values.

The package constructs six cases:

Closure	Routes	Modes	Congestion
`H`: Hulten	Not solved	Not solved	Traditional approach
`NC`: no congestion	Flexible	Flexible	None
`NT`: no terminals	Flexible	Flexible	Edge congestion
`F`: full	Flexible	Flexible	All specified congestion
`FM`: fixed modes	Flexible	Fixed at baseline shares	All specified congestion
`FR`: fixed routes	Baseline OD-edge incidence fixed	Flexible	All specified congestion

The labels describe local derivatives, not finite counterfactual models. In particular, `FR` holds the differentiated route incidence fixed while retaining modal response and the remaining equilibrium system. Both spatial models use these comparisons. The additional normalized and allocation–scarcity–equilibrium decompositions below are available only for the economic-geography model.

4.3.2 Economic-geography normalized multipliers

To compare equilibrium propagation across cases, divide out traffic and the common scale ρ. For a positive-traffic edge e with $\rho \neq 0$, define

$$m_{e,S} = \frac{E^r_{e,S}}{\rho \Xi_e}. \tag{4.8}$$

The resulting multiplier measures equilibrium propagation per unit of traffic. Define the signed wedges

$$d_{edge} = m_{NC} - m_{NT}, \quad d_{terminal} = m_{NT} - m_F,$$
$$d_{mode} = m_F - m_{FM}, \qquad d_{route} = m_F - m_{FR}.$$

Evaluating them at the common baseline yields the closure ladder

$$m_F = m_{NC} - d_{edge} - d_{terminal} = m_{FM} + d_{mode} = m_{FR} + d_{route}. \tag{4.9}$$

The terms are signed. A positive road wedge in multiplier units need not imply a positive contribution to the welfare elasticity when $\rho < 0$.

The first equality gives an additive congestion ladder. The other two compare the full model separately with fixed modes and fixed routes. Mode and route wedges are not added to the congestion ladder because each is an alternative closure comparison.

The urban model reports the corresponding differences directly in welfare elasticity levels:

$$E^r_{U,F} - E^r_{U,NC}, \qquad E^r_{U,F} - E^r_{U,FM}, \qquad E^r_{U,F} - E^r_{U,FR}.$$

It does not divide by the economic-geography coefficient ρ. The same comparisons therefore apply without imposing a trade-model normalization on residence–workplace welfare.

4.3.3 Inverse-gap identity

Each rung compares two derivatives that share the welfare row and policy forcing but use different Jacobians. The matrix identity $J_A^{-1} - J_B^{-1} = J_A^{-1}(J_B - J_A)J_B^{-1}$ therefore implies

$$q_c^\top J_{c,A}^{-1} b_c - q_c^\top J_{c,B}^{-1} b_c = q_c^\top J_{c,A}^{-1}(J_{c,B} - J_{c,A}) J_{c,B}^{-1} b_c.$$

The package checks this identity for the edge, terminal, mode, and route comparisons in both spatial models. It is a direct algebraic test of the signs, orientation, and Jacobian construction.

4.3.4 Economic-geography analytical components

The inverse-gap identity gives the total difference between two closures. For economic geography, the package further separates each mode or route wedge into allocation, congestion-scarcity, and aggregate-equilibrium components:

$$d_r = d_r^{alloc} + d_r^{scar} + d_r^{eq}, \qquad r \in \{mode, route\}.$$

Under the common-baseline construction used here, the sparse spatial allocation block and aggregate-feedback term are the same across the five closures. Their contributions are therefore exact zeros for the road, terminal, mode, and route comparisons, and the congestion-scarcity component carries each wedge. The package nevertheless evaluates the general mixed update for the mode and route comparisons and checks both the implied zeros and the reconstruction of the total wedge. Appendix B gives the Jacobian blocks, rank-two update, and structured inverse formulas.

4.3.5 The economic-geography Hulten gap

Substituting $E^r_{e,F} = \rho \Xi_e m_{e,F}$ from (4.8) into the difference between traffic and the full realized derivative gives

$$\Xi_e - E^r_{e,F} = \Xi_e(1 - \rho) + \rho \Xi_e (1 - m_{e,F}).$$

The first term is the common externality scaling. The second is equilibrium propagation relative to the Traditional approach.

For a nonzero realized derivative, define $\chi_e^{eff} = E^\theta_{e,F} / E^r_{e,F}$. Next substitute the first equality in (4.9) and use $E^{\theta}_{e,F} = \chi^{eff}_e E^r_{e,F}$. This produces

$$\Xi_e - E^{\theta}_{e,F} = \Xi_e(1-\rho) + \rho\Xi_e(1-m_{e,NC}) \\ + \rho\Xi_e d_{edge,e} + \rho\Xi_e d_{terminal,e} + \rho\Xi_e m_{e,F}(1-\chi^{eff}_e).$$

A zero realized derivative leads the output to compute the primitive-minus-realized contribution directly in levels instead of forming a ratio. Signed percentages can be misleading if large positive and negative components nearly cancel.

The identity reports all terms in units of the Traditional approach minus primitive-cost welfare. A positive contribution widens that gap; a negative contribution offsets it. The sign convention applies even when $\rho < 0$.

For the urban model, the package reports $\Xi_e - E^r_{U,F}$ and $\Xi_e - E^{\theta}_{U,F}$, together with the closure gaps above. It does not label these differences as externality scaling or trade-market-access propagation because those labels come from the economic-geography factorization.

4.3.6 Physical-link aggregation

The model is directed. To obtain a two-direction physical-link experiment, the software first sums the primitive and realized elasticities across the two directions. It applies the same operation to Hulten traffic, all additive Hulten-gap components, and all analytical components. Only after summation does it calculate a physical-link normalized multiplier. This order avoids averaging ratios with different traffic denominators.

> **Practical caution.** A physical-link result is defined only when exactly two declared policy arcs have opposite endpoints and a common physical-link identifier. Transit services and one-way links that lack such a pair remain directed policies.

4.3.7 Reading a decomposition

Begin with the gap between the Traditional approach and the full primitive-cost result. In the economic-geography model, separate the common externality scale from link-specific equilibrium propagation before examining edge and terminal congestion. Next compare the primitive- and realized-cost effects to measure pass-through. The fixed-mode and fixed-route cases then show whether modal or route adjustment contributes to ranking changes. Interpret these economic differences only after the identity residuals have passed their numerical checks. The decomposition explains a result within the specified model. It does not show which parameters or model assumptions are empirically appropriate.

Applying these identities requires a baseline whose flows, units, and policy definition satisfy the model's accounting restrictions. Subsection 5.1 states that empirical contract before the guide turns to the software.

Part III Detailed guide: implementation and applications

The derivative in Part II is a theoretical result, not yet an empirical estimate. Part III explains how data, calibration, and software produce that object in an application. It then verifies the calculation on a public grid and relates the same workflow to the U.S. freight and urban-commuting applications.

5 Implementation and verification

Part II identifies the traffic, equilibrium, and policy objects required by the welfare derivative. An application must represent those objects in common units, construct a coherent baseline, and preserve the policy definition through the numerical calculation. Data and calibration come first because the validity of later software checks depends on that representation.

5.1 Data and calibration

Raw transport records rarely coincide with the model's nodes, flows, or cost units. This subsection states the model-ready schemas, accounting identities, parameter sources, and policy fields that convert those records into inputs the derivative can use.

5.1.1 Model-ready inputs

The theory identifies the inputs needed by the derivative; the data contract states how an application must represent them. The generic adapter reads two CSV files and one TOML configuration. The files are deliberately narrow. They contain a baseline that already satisfies the selected model's accounting identities; they are not a warehouse for raw transport records.

For economic geography, the node file contains:

```
node_id,labor,income,longitude,latitude
```

Labor and income shares must be finite, strictly positive, and normalized. Zero-mass locations are outside the interior baseline accepted by the generic loader. Coordinates are optional for analysis but required for maps.

For urban commuting, the node file instead contains:

```
node_id,residents,employment,longitude,latitude
```

Residence and workplace margins must be finite, positive, and normalized. Their geographic distributions may differ even though their aggregate totals refer to the same commuter population.

The edge-mode file contains one row for each active mode on each directed edge:

```
edge_id,physical_link_id,origin,destination,mode,flow,
origin_terminal_id,destination_terminal_id,
congestion_elasticity
```

The terminal and edge-specific congestion columns are optional unless selected by the configuration. The pair (`edge_id`, `mode`) must be unique. One directed `edge_id` may therefore appear once for each active mode, but all such rows must share endpoints and a physical-link identifier. Each endpoint must exist in the node table.

Implementation. Recovering baseline cost derivatives from flows. Given the recursive-route and modal-logsum assumptions, consistent edge–mode flows and node margins reveal the continuation and modal shares needed for the local cost derivative. The package reconstructs those shares without separately observed route-level generalized costs. This identification also requires calibrated route and modal elasticities and flows that satisfy the model's accounting identities. Subsection 4.2.7 distinguishes observed, reconstructed, calibrated, and computed quantities.

5.1.2 Flow concepts and units

Economic geography uses value traffic normalized by world income. The urban model uses commuter traffic normalized by total commuters. Vehicle counts, passenger trips, tons, and scheduled services are not automatically interchangeable with either flow concept. A preprocessing pipeline may convert a physical measure to the model's flow, but that conversion must be explicit and defensible.

All active modal flows on an edge sum to aggregate edge traffic:

$$\Xi_{ij} = \sum_m \Xi_{ij,m}, \qquad s_{ij,m} = \frac{\Xi_{ij,m}}{\Xi_{ij}}.$$

The package constructs modal shares from the input baseline. It does not use η to overwrite those baseline shares. The elasticity controls the local response after a shock.

For economic geography, labor and income shares each sum to one, and edge value flows use the same world-income denominator. For urban commuting, residence and workplace margins each sum to one, and edge flows use the same total-commuter denominator. In either model, modal shares sum to one within a directed edge. Edge traffic does not generally sum to one because the same shipment or commuter can traverse several edges along a route. Renormalizing the edge table to unit total would therefore change the welfare scale and violate the route accounting.

The unit contract extends to congestion. A log elasticity such as $\gamma = 0.092$ is invariant to proportional rescaling of a correctly matched flow measure. It is not invariant to replacing value or commuter traffic with an unrelated vehicle, passenger, or schedule measure.

5.1.3 Accounting checks

For economic geography, outgoing and incoming traffic must recover the same trade exposure:

$$\mathcal{T}_i^G = \frac{O_i}{1 - s_i^x} = \frac{I_i}{1 - s_i^y}.$$

For urban commuting, residence and workplace absorption must recover the same commuting exposure:

$$\mathcal{T}_i^U = l_i^F + O_i = l_i^R + I_i.$$

Using the identity selected by the configuration, the loader first verifies node and mode identifiers, positivity, and finiteness. It then checks flow conservation and agreement between alternative measures of the exposure stock. For the requested policy, it also verifies reciprocal physical-link pairing, coverage of the policy mode, terminal identifiers for modes subject to terminal congestion, and congestion keys against the observed active modes.

Perform any symmetrization, padding, modal rescaling, or balancing before loading the generic adapter. When empirical sources do not agree, create a preprocessing artifact that states the transformation and supply its output to the package. The run manifest records the hashes of the transformed inputs.

5.1.4 Preparing model-ready flows

Different sources require different preprocessing. This table identifies the main empirical gap without presuming a universal conversion.

Available information	Required construction
Directed edge value flows	Apply unit, price-basis, and accounting checks directly.
Vehicle or passenger counts	Map the physical count to the model's value or commuter flow and match the congestion elasticity to that measure.
Origin–destination value flows	Assign each OD flow to directed edge–modes under a declared routing model.
Scheduled transit service	Combine schedules with ridership or a modeled commuter baseline; schedules alone are not observed traffic.
Imports and exports with different margins	Project them onto declared common margins while preserving and reporting the chosen support.
Incomplete modal coverage	Define a defensible residual mode or limit the application to the observed choice set.
Zero-flow proposed infrastructure	Use a finite counterfactual or entry model; the interior local derivative does not identify the new alternative.

Reconcile inconsistent source matrices in a separate, versioned balancing step. First choose the target node margins and preserve structural zeros that represent unavailable connections. Then solve a constrained projection, such as iterative proportional fitting when its support conditions hold, that changes the source matrix as little as possible while matching those margins. Report the projection rule, fixed margins, support changes, convergence tolerance, iteration count, maximum margin error, and a distance between raw and balanced flows. Archive both the raw-source hashes and the model-ready output hash. The freight adapter below implements this procedure for directional port trade; other applications must justify their own targets and support.

5.1.5 Parameters

After the accounting inputs are fixed, the configuration supplies the elasticities governing spatial, route, modal, and congestion responses. The economic-geography parameter vector is

$$(\alpha, \beta, \sigma, \eta, \gamma_{e,m}).$$

Here $\gamma_{e,m}$ denotes the declared edge–mode congestion elasticity; the route curvature is tied to $\sigma - 1$ by the paper's theorem. The paper baseline uses

$$\alpha = 0.10, \quad \beta = -0.30, \quad \sigma = 9, \quad \eta = 1.099, \quad \gamma_{\text{road}} = 0.092.$$

These values are external calibration inputs, not parameters estimated by the package. The spatial elasticities $\alpha = 0.10$ and $\beta = -0.30$ follow the economic-geography calibration in Allen and Arkolakis (2014), as used in the December 2020 revision of Allen and Arkolakis (2020). The route curvature $\sigma - 1 = 8$ matches the value used by Fuchs and Wong (2026), while their estimated modal elasticity gives $\eta = 1.099$. The road-congestion elasticity $\gamma_{\text{road}} = 0.092$ comes from Allen and Arkolakis (2020). The endpoint-terminal extension uses the 0.096 dwell-time elasticity estimated by Fuchs and Wong (2026). These citations identify the source of each baseline value; they do not substitute for sensitivity analysis in a new application.

The urban parameter vector replaces σ with commuting heterogeneity θ_U:

$$(\alpha, \beta, \theta_U, \eta, \gamma_{e,m}).$$

The package uses θ_U as the urban route curvature. Its modal and congestion elasticities remain application inputs. The Seattle candidate transfers $\eta = 1.099$ from the freight application and reports sensitivity rather than presenting it as a Seattle estimate.

Sensitivity should vary one parameter while remaining on a regular branch. The paper sensitivity runner maintains

$$\alpha + \beta < 0, \qquad e < 0, \qquad |e| \geq 0.05, \qquad 1 + \alpha + \beta > 0.05.$$

Thus “net dispersion” means $\alpha + \beta < 0$, and the remaining conditions keep the path on the baseline sign branch and away from the two transformation singularities. The generic model builder requires the corresponding denominators to be nonzero. Both models reject singular route, transport, or spatial systems, and the economic-geography implementation rejects an independent route curvature not covered by the theorem.

5.1.6 Policy definition

The configuration states the policy mode and whether the output concerns directed arcs, reciprocal physical links, or both. It also records the proportional shock used to translate elasticities into reported gains and indicates whether the policy combines several edge–mode shocks in a bundle.

A policy bundle is a sparse list:

```
bundle_id,edge_id,mode,weight
```

The first-order welfare effect is linear in these weights. Duplicate entries, unknown edges, inactive modes, and nonfinite weights are rejected.

5.1.7 External-data adapters

The freight replication adapter, retained in code under its historical `RSUE` name, performs transformations required to reproduce the paper’s research pipeline. Unlike the generic adapter, it may pad, symmetrize, filter, or rescale because those operations are part of the declared replication contract. Each transformation is listed in TOML and written to the manifest.

The Seattle adapter combines a published urban archive with LODES commuting, ACS mode shares, and a pinned historical GTFS network (U.S. Census Bureau, 2026b, 2017; King County Metro, 2017a). The raw sources remain outside Git. Schedule data identify service paths and frequency, not observed route ridership. The model builder therefore reports its construction and external comparisons without presenting schedule counts as passenger flows.

> **Practical caution.** A successful schema check establishes internal consistency, not external validity. Validation against untargeted traffic, ridership, prices, or location outcomes is a separate empirical exercise.

Once the data and policy satisfy this contract, the computation constructs the route basis, assembles the relevant equilibrium system, and converts each policy forcing into a welfare effect. The next subsection follows that sequence.

5.2 Software workflow and outputs

Once the inputs satisfy the data contract, the software must carry them through the same sequence as the derivation: construct the transport block, form the spatial Jacobian, solve the adjoint, and write results with their diagnostics and provenance. The public interface exposes that sequence without requiring an application to modify the solver.

5.2.1 Public interface

The package exposes the data-to-derivative sequence through a small typed API. The same calls validate inputs, build the model selected in the configuration, and write results with their provenance:

```
using TransportNetworkWelfare

project = load_project("examples/grid_multimodal/config.toml")
validate(project)
model = build_model(project)
welfare = welfare_effects(model)
components = decompose_welfare(model)
sensitivity = sensitivity_path(model, :eta, [0.9, 1.099, 1.25])
write_results(components, "examples/grid_multimodal/output/guide-api/"; project)
```

The configuration determines the spatial, modal, congestion, and policy types. Users adapting a dataset should change CSV and TOML inputs rather than editing the solver.

The CLI calls the same public functions:

```
julia --project=. bin/tnw.jl validate \
  examples/grid_multimodal/config.toml
julia --project=. bin/tnw.jl analyze \
  examples/grid_multimodal/config.toml
julia --project=. bin/tnw.jl decompose \
  examples/grid_multimodal/config.toml
julia --project=. bin/tnw.jl sensitivity \
  examples/grid_multimodal/config.toml
```

5.2.2 Project isolation and package updates

Section 1 gives the initial setup and run commands. Keep the resulting empirical project outside the package checkout so application inputs and outputs remain separate from package updates. Before updating, preserve the package commit and manifest used for archived results. Then fast-forward the package and rerun its tests:

```
git rev-parse HEAD
git switch main
git pull --ff-only
julia --project=. -e 'using Pkg; Pkg.instantiate(); Pkg.test()'
```

Revalidate and rerun the external project after any package, configuration, or data change. A changed output belongs to a new result vintage; retain the earlier commit, inputs, and manifest when earlier results must remain reproducible.

5.2.3 Implementation boundary

Users normally interact with the API or CLI rather than individual source modules. The implementation follows the derivation's dependency order: validate data, reconstruct routes and modes, construct the selected spatial system, eliminate congestion feedback, and project each policy onto welfare. Appendices A and B map these operations to source files and functions. The tracked theory–code audit records the line-level correspondence for contributors.

5.2.4 Build stages

Figure 1 in the quick start summarizes the calculation. Each arrow passes a checked result to the next stage, beginning with model-ready inputs and ending with archived results.

For either spatial model, `build_model` begins by loading and validating the project and network, identifying the positive-flow edge–mode pairs, and reconstructing the route kernel and bilateral flow matrix. It uses these objects to form the flexible and fixed route derivatives. The economic-geography or urban builder then supplies the no-congestion spatial Jacobian, cost loading, and welfare row.

The remaining stages construct the modal and congestion maps for each decomposition case and eliminate congestion quantities through the Schur complement. After checking the route, transport, and equilibrium condition numbers, the builder exposes the equilibrium system used for the welfare and decomposition calculations.

`analyze` uses a welfare-only builder. It constructs the full flexible case but omits the dense fixed-route incidence matrix. `decompose` constructs all decomposition cases. Omitting fixed-route incidence from ordinary analysis reduces memory use because that matrix scales with origins, destinations, and edges.

5.2.5 Policy derivatives

For all directed policy arcs, the code constructs a sparse selector matrix. The realized forcing changes the aggregate cost on the selected edge. The primitive forcing first passes the edge–mode shock through the modal and route responses. It then incorporates edge or terminal congestion, aggregates the result to the directed edge, and loads the resulting cost change into the spatial residual.

The adjoint gain function solves the transposed Jacobian and multiplies the forcing matrix. It returns one elasticity per policy arc. The decomposition then applies inverse-gap and structured-update identities to the same forcing.

5.2.6 Outputs and manifests

Analysis writes directed results to `welfare_directed.csv` and, when requested, reciprocal-link results to `welfare_physical.csv`. The `decompose` command adds the corresponding decomposition tables. Each run also writes `diagnostics.json` and `run_manifest.json`.

The manifest records the package version and Git state, Julia and BLAS environment, configuration hash, input hashes, model variants, parameters, transformations, output hashes, and verification status. JSON serialization is strict: unsupported values and colliding stringified keys are errors rather than silently converted strings.

Field	Meaning	Unit or sign
`physical_link_id`	Identifier shared by the reciprocal directed policies.	Identifier.
`hulten`	Traditional approach: baseline policy-mode value traffic.	Positive welfare elasticity.
`realized_F`	Full-model response to a reduction in the cost users face.	Benefit elasticity; may be zero.
`primitive_F`	Extended approach: full-model response to the primitive infrastructure cost.	Headline benefit elasticity.
`chi_effective`	Ratio `primitive_F/realized_F`.	Missing when the denominator is numerically zero.
`primitive_externality` `primitive_propagation`	Common externality scale and link-specific equilibrium-propagation terms in the signed Hulten gap.	Signed welfare elasticities.
`primitive_edge` `primitive_terminal`		

Field	Meaning	Unit or sign
`primitive_pass_through`	Edge congestion, terminal congestion, and primitive-to-realized cost contributions.	Signed welfare elasticities.
`identity_residual_*` and `channel_residual_*`	Reconstruction errors for the closure ladder and analytical-component identities.	Absolute numerical residuals; target below 10^{-10}.

The `decomposition_physical.csv` file sums opposite-direction elasticities and signed components before calculating link-level ratios and normalized multipliers. Ranks are derived from the selected welfare column; they are not stored as primitive output fields. Join the physical-link file to engineering costs using `physical_link_id`, and preserve the mapping used for that merge. A link is ready for economic comparison only when the run-level accounting, conditioning, and identity diagnostics pass.

5.2.7 Computational cost

The sparse welfare calculation is appropriate for many candidate edges on a moderate network. The exact route decomposition is more demanding because it constructs baseline origin-destination edge incidence. Before decomposition, the validator reports an incidence-memory estimate. Large traffic-assignment standard networks may be suitable for the welfare derivative but not the full fixed-route decomposition on a laptop.

The package does not replace an ill-conditioned model with low-rank fitting, ridge regularization, or dropped routes. Larger applications should exploit sparsity or use a separately derived route comparison while retaining the verification gates.

The public grid exposes the intermediate quantities and diagnostics used in the calculation. Before applying the architecture to that baseline, the next subsection states the numerical and reporting checks that a calculation must pass.

5.3 Numerical checks and reporting

A numerical solve is not sufficient evidence that a reported elasticity represents the intended model. The inputs, derivative identities, and empirical mapping require separate checks. Distinguishing these layers also prevents a successful implementation test from being described as external validation.

5.3.1 Three layers of verification

A credible application must pass accounting, mathematical, and empirical tests. The accounting test asks whether the input margins, exposure stocks, routes, and edge flows describe the same baseline. The mathematical test asks whether the code implements the declared derivative and closure identities. The empirical test asks whether the calibrated model can explain outcomes that were not imposed in constructing that baseline. The package automates the accounting and mathematical tests; external validation requires application-specific evidence.

5.3.2 Route diagnostics

The run manifest records the route spectral radius and the absorption error. It also reports the bilateral row and column errors and the directed-edge reconstruction errors under flexible and fixed routing. The spectral radius must be below one. A value close to one can produce a well-defined but poorly conditioned resolvent in which traffic circulates for many expected steps before absorption.

Large margin or edge errors usually indicate inconsistent inputs, reversed orientation, duplicated arcs, or an undeclared transformation. They should not be repaired by increasing the tolerance.

5.3.3 Linear-system diagnostics

Each decomposition case reports the condition numbers of its transport and spatial systems. The default upper gate is 10^{12}. This is a permissive numerical ceiling rather than a target. Results near the gate can be sensitive to small changes in data or parameters.

The package also records solve residuals and finiteness checks. A low residual with a very high condition number does not establish a stable economic derivative. Both should be examined.

5.3.4 Identity diagnostics

For `decompose`, both spatial models report the maximum inverse-gap and closure-ladder errors, the realized and primitive Hulten-accounting errors, the closure-Jacobian reconstruction error, and the primitive transport residual. Economic-geography output also reports a reconstruction error for its three analytical subchannels: allocation, scarcity, and equilibrium correction. That subcomponent decomposition is not defined for the urban model. The package acceptance target for algebraic identities is 10^{-10}. An identity failure is a code or model-construction problem. It is not an interesting economic result.

5.3.5 Finite differences

An analytical derivative should be compared with independently solved nonlinear equilibria. Let h be a small change in primitive log cost and re-solve welfare on either side of the baseline. The central-difference estimate is

$$E_p^{FD} = -\frac{\log W(\theta_p + h) - \log W(\theta_p - h)}{2h}.$$

The nonlinear solver must re-solve all equilibrium and transport quantities. Reusing the analytical derivative inside the counterfactual would make the test circular.

The package tests efficient, externality, congestion, asymmetric, multiple mode, unique-route, and parallel-route cases. The archived freight verification uses a deterministic traffic-stratified sample of road arcs. The absolute derivative error must be below 10^{-6}, and nonlinear solve residuals below 10^{-10}.

5.3.6 First-order accuracy

Finite-difference verification establishes the derivative at the baseline. To assess a finite policy, compare the first-order prediction

$$\Delta \log W_{\text{linear}} = -E_p \Delta\theta_p$$

with a nonlinear counterfactual. The package checks changes of 0.1, 1, and 5 percent only for models with an independently implemented nonlinear solver. If the derivative and nonlinear solver are correct and the equilibrium is smooth, the approximation error should shrink at the rate implied by a first-order expansion as the shock approaches zero.

The urban module supplies such a nonlinear solver. The current economic-geography grid verifies the local derivative against a re-solved nonlinear route–mode–congestion system, but it does not supply a global economic-geography counterfactual for a five-percent project. A finite-size accuracy chart for that model would therefore overstate what has been verified. Analysts should report this distinction and the policy scale used with the approximation.

5.3.7 Sensitivity

A sensitivity path should report both average effects and rank stability. A parameter can scale all links similarly while leaving the ordering nearly unchanged, or leave the average stable while reshuffling the top links.

The paper exercises vary α, β, net dispersion, η, and road congestion in the headline model. Common road/terminal and terminal-only paths are extension diagnostics. Each point must remain on the declared regular branch and pass all numerical gates.

For an urban application, sensitivity replaces the trade elasticity with commuting heterogeneity θ_U

and should separately vary the modal and congestion inputs used by that application. A transferred value of η or an unestimated terminal elasticity should be labeled as such.

At each accepted point on a sensitivity path, report the mean and median primitive-cost welfare elasticities together with their Pearson and Spearman correlations with the baseline. Ranking stability should be summarized by top-k overlap, maximum absolute rank movement, and the number of links that enter or leave the top decile. The report should also identify sign changes, changes in the signed decomposition, transport and spatial condition numbers, and any branch failures or rejected parameter values. One-at-a-time paths isolate separate parameter effects. Joint low, central, and high scenarios are worth reporting when spatial externalities, congestion, or modal substitution are expected to move together. These exercises report rankings for each assumed calibration; the package does not estimate sampling uncertainty.

5.3.8 Common failures

Message or symptom	Likely cause	Correct response
Route kernel is not contractive	Cyclic continuation shares do not absorb	Revisit flow accounting or local absorption; do not truncate paths silently.
Stock disagreement	Incoming and outgoing flows imply different exposure	Document and perform a balancing transformation outside the package.
Policy edge lacks active mode	Zero or missing policy-mode flow	Remove the policy or provide a defensible positive-flow baseline.
Missing terminal ID	Positive terminal congestion on an unidentified terminal	Add explicit endpoint terminal identifiers.
Condition-number gate	Near-singular modal, transport, or spatial system	Review parameters, baseline support, and model identification.
Undefined effective pass-through	Realized derivative is numerically zero	Use the level pass-through component; do not force a ratio.
Physical link rejected	Missing or inconsistent reciprocal direction	Report directed policies or correct the physical-link identifier.

5.3.9 Reporting standards

An empirical application should disclose how the network and flows were constructed, the units in which they are measured, and all transformations used to make them model-ready. It should define the policy primitive, direction, mode, and scale and report the sources of all parameters together with the modal and congestion specifications.

The results should be accompanied by the route, accounting, conditioning, and identity diagnostics, as well as sensitivity of both levels and rankings. Finally, the application should distinguish external-validation evidence from quantities matched by construction and state the limitations of that evidence. This information is part of the economic result, not software housekeeping.

> **Implementation. Copy-ready reporting template.** "We evaluate a s-percent reduction in the primitive generalized cost of each existing [mode] link. The headline statistic is the full primitive-cost welfare elasticity. We compare it with the Traditional approach and the full realized-cost effect. We report the signed congestion, equilibrium-propagation, and pass-through components; Pearson and Spearman correlations; top-k overlap and rank movements; route and accounting errors; condition numbers; finite-difference errors; and parameter sensitivity. The statistic measures model-implied gross economic benefits and excludes construction costs and omitted externalities."

The public grid in Section 6 satisfies these numerical checks by construction. Section 7 then separates that generic verification from the additional data and validation claims required by the U.S. freight and Seattle applications.

6 Worked example: a multimodal grid

The theory, data contract, and software workflow have so far been presented separately. The grid reunites them in one reproducible calculation. Its small network makes the route and modal mechanisms visible while exercising the same validation, welfare, decomposition, and plotting commands used by larger applications.

6.1 Economic setting

The worked example uses the economic-geography model. Transit is a second mode in its transport network; the example does not use the urban residence–workplace model. It has 25 economic locations arranged on a five-by-five grid. Adjacent locations are connected by reciprocal road arcs. Transit is also available on the four links along the central east–west axis and has a larger baseline modal share there. The committed mode shifters rationalize these baseline shares; η governs their response to a cost change rather than identifying relative cost levels from shares alone. Each transit stop has an explicit terminal identifier, so the example can include congestion from traffic using the two endpoint terminals of a trip. Figure 3 shows the network.

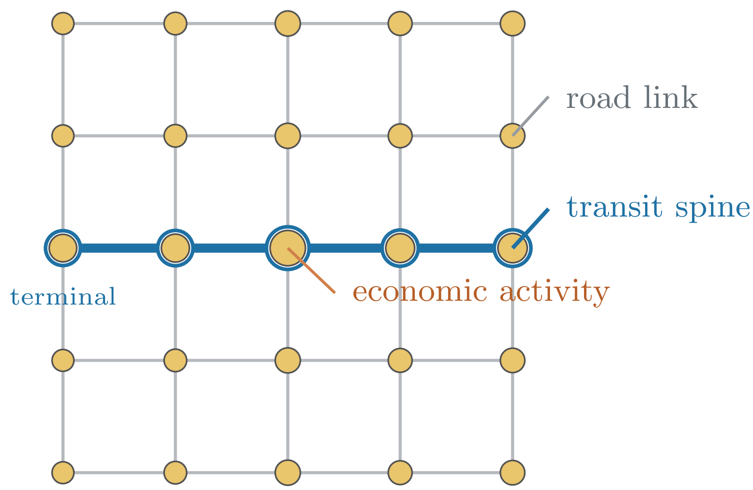

Figure 3: Synthetic road grid and transit spine. Thin gray links are roads; thick blue links also carry transit. Gold node area represents economic activity, and blue rings identify transit terminals. Road and transit connections are reciprocal.

Activity is positive throughout the grid and is higher near the central spine, the north–south cross-axis, and the central node. A small eastward gradient breaks otherwise arbitrary ranking ties. The input builder assigns symmetric directional flows and generates location margins from the same accounting system. The package therefore reconstructs the route kernel without silently balancing the network.

To keep the equilibrium forces comparable with the paper while making the network fully reproducible, the example uses

$$\alpha = 0.10, \qquad \beta = -0.30, \qquad \sigma = 9, \qquad \eta = 1.099.$$

These reference values mirror the paper calibration described in Section 5.1, drawing on Allen and Arkolakis (2014, 2020) and Fuchs and Wong (2026). Modes enter through `ChoiceLogsum`. Road congestion has elasticity 0.08, and the illustrative endpoint-terminal elasticity for transit is 0.025. The two congestion values make the model comparisons visible; they are not estimates for a particular city.

6.2 Build the inputs, run the model, and draw the results

The committed input has 80 directed road arcs, 40 reciprocal road links, and 88 positive-flow edge–mode pairs. The difference between the first and third counts is the eight directed transit edge–mode rows on the central spine.

The input builder is `examples/grid_multimodal/build_inputs.jl`. It first assigns positive labor and income across the grid. It then writes reciprocal horizontal and vertical road arcs and adds transit rows only along the central axis. The following excerpt shows the construction that an empirical preprocessing script must replace:

```
const GRID_SIZE = 5
const TRANSIT_ROW = 3

node_id(row, column) = @sprintf("r%dc%d", row, column)

for row in 1:GRID_SIZE, column in 1:(GRID_SIZE - 1)
    west = node_id(row, column)
    east = node_id(row, column + 1)
    link = @sprintf("H_r%d_c%d", row, column)
    road = road_flow(:horizontal, row, column)
    edge_row(io, link * "_E", link, west, east, "road", road)
    edge_row(io, link * "_W", link, east, west, "road", road)
    if row == TRANSIT_ROW
        transit = transit_flow(column)
        edge_row(io, link * "_E", link, west, east,
                 "transit", transit;
                 origin_terminal=terminal_id(row, column),
                 destination_terminal=terminal_id(row, column + 1))
        edge_row(io, link * "_W", link, east, west,
                 "transit", transit;
                 origin_terminal=terminal_id(row, column + 1),
                 destination_terminal=terminal_id(row, column))
    end
end
```

The complete script also writes vertical links and explicit origin and destination terminal identifiers. Symmetric directional flows make incoming and outgoing traffic equal at each node. This is a transparent device for a synthetic example, not a balancing rule that should be imposed on empirical commuting or freight data.

Rebuild the CSVs and inspect the first rows before running the model:

```
julia examples/grid_multimodal/build_inputs.jl
head examples/grid_multimodal/data/nodes.csv
head examples/grid_multimodal/data/edge_modes.csv
```

The configuration names the input files, selects `economic_geography` and `ChoiceLogsum`, declares road and terminal congestion, and defines a one-percent bidirectional road policy. These are the fields an analyst changes after replacing the CSVs:

```
[model]
spatial_specification = "economic_geography"
alpha = 0.10
beta = -0.30
sigma = 9.0
eta = 1.099
```

```
modal_specification = "choice_logsum"

[policy]
mode = "road"
unit = "both"
shock_fraction = 0.01
```

Validation should precede the numerical work. From the repository root, run:

```
julia --project=. bin/tnw.jl validate \
  examples/grid_multimodal/config.toml
julia --project=. bin/tnw.jl analyze \
  examples/grid_multimodal/config.toml
julia --project=. bin/tnw.jl decompose \
  examples/grid_multimodal/config.toml
julia --project=. bin/tnw.jl sensitivity \
  examples/grid_multimodal/config.toml
```

The analysis writes directed and physical-link welfare tables. The decomposition adds the closure and component columns used below, while the sensitivity command writes one row for each declared parameter value. Each run also updates `diagnostics.json` and `run_manifest.json`.

The figures, tables, and TeX macros in this section are generated from those validated outputs rather than transcribed by hand. The documentation asset target rebuilds the inputs, reruns the example, and redraws the figures:

```
make practitioner-guide-assets
```

For an external application, the equivalent standalone map command is given in Subsection 1.6; it reads `welfare_physical.csv` and can use line and polygon GeoJSON layers.

The policy improves one physical road link at a time by lowering primitive cost by one percent in both directions. Transit remains available throughout the equilibrium response. The result therefore measures how a road improvement affects a network in which users can also change mode.

6.3 Baseline checks

The reconstructed route kernel has spectral radius 0.369, safely below one. The full equilibrium Jacobian has condition number 9.16×10^3, below the package gate of 10^{12}. The maximum inverse-gap residual is 1.10×10^{-15}, the closure-ladder residual is 0, and the largest analytical component-reconstruction residual is 1.65×10^{-13}. These checks establish that the reported derivatives solve the stated numerical model. They do not turn the synthetic parameters into empirical estimates.

6.4 Traffic and primitive-cost rankings

The mean Traditional-approach elasticity is 0.010520, compared with a mean primitive-cost welfare elasticity of 0.005425. The median and maximum primitive-cost elasticities are 0.004871 and 0.010177. For the configured 1.0 percent primitive improvement, the mean primitive-cost elasticity implies a first-order log-welfare gain of 0.00005425, approximately 0.005425 percent. This is a local benefit prediction, not a project net present value.

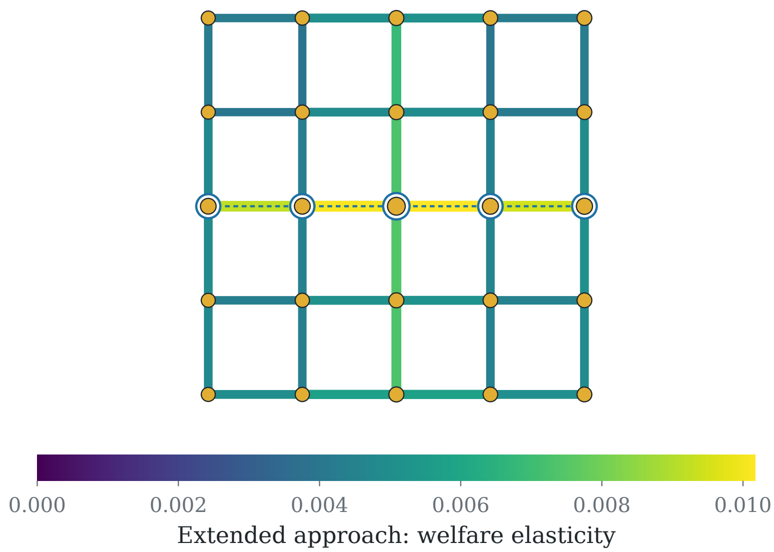

Figure 4: Extended-approach primitive-cost welfare elasticities for physical road-link policies. Multiplying an elasticity by 0.01 gives the first-order welfare gain from the configured one-percent improvement. The blue line and terminal rings mark links on which transit is also available.

Benefits are largest on the central row and column, where economic activity and network use are concentrated. The transit spine does not make the central road links unimportant. It changes how their direct cost savings propagate through modal allocation and congestion.

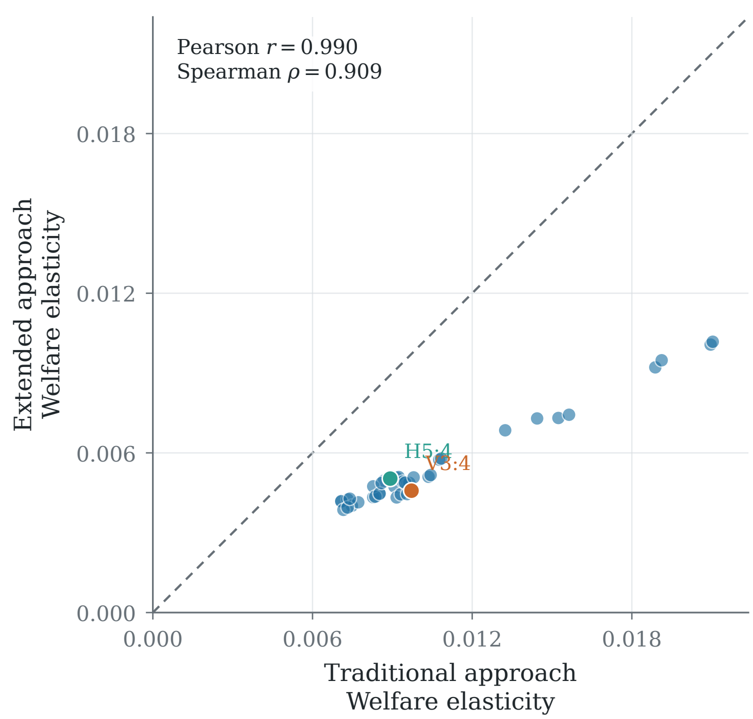

Figure 5: Traditional- and Extended-approach welfare elasticities for the 40 physical road improvements. The dashed line is the 45-degree line. Labels identify the largest upward and downward rank changes.

The Pearson correlation is 0.990, and the Spearman correlation is 0.909. The close linear relation reflects the importance of baseline traffic. It does not imply identical rankings. Link H_r5_c1 rises from 26 under the Traditional approach to 17 under the primitive-cost effect. Link V_r3_c2 falls from 17 to 28. These links have similar traffic but differ in their access to the high-activity cross-axis and the transit spine. The Extended approach values that difference in equilibrium incidence.

Table 1: Highest-ranked physical road links under the primitive-cost welfare effect

Link	Traditional	Extended	Traditional rank	Extended rank
Row 3, columns 3-4	0.021040	0.010177	1	1
Row 3, columns 2-3	0.020960	0.010069	2	2
Row 3, columns 4-5	0.019120	0.009481	3	3
Row 3, columns 1-2	0.018880	0.009214	4	4
Column 3, rows 3-4	0.015640	0.007433	5	5
Column 3, rows 2-3	0.015240	0.007312	6	6
Column 3, rows 4-5	0.014440	0.007294	7	7
Column 3, rows 1-2	0.013240	0.006849	8	8

The highest-ranked link, Row 3, columns 3–4, has a Traditional-approach elasticity of 0.021040, a full realized-cost effect of 0.022809, and a primitive-cost effect of 0.010177. The highest-ranked links are stable across the traffic and primitive-cost measures. Ranking changes occur farther down the list, where direct traffic differences are smaller and equilibrium propagation has more scope to matter. A high correlation or agreement at the top therefore does not establish that the two criteria are equivalent.

6.5 Closure ladder

The common-baseline ladder separates the modeled effects without recalibrating the starting equilibrium. Applying the cases defined in Section 4.3, the mean realized-cost elasticities are

	NC	NT	F	FM	FR
mean elasticity	0.010594	0.012138	0.012279	0.012280	0.012621.

Moving from NC to NT adds road congestion; moving from NT to F adds terminal congestion. Comparing F with FM or FR isolates the local mode and route wedges. Their mean normalized values are 0.000050 and 0.020109, respectively. Because these are normalized multipliers, their signs should be interpreted together with ρ, not as welfare levels on their own.

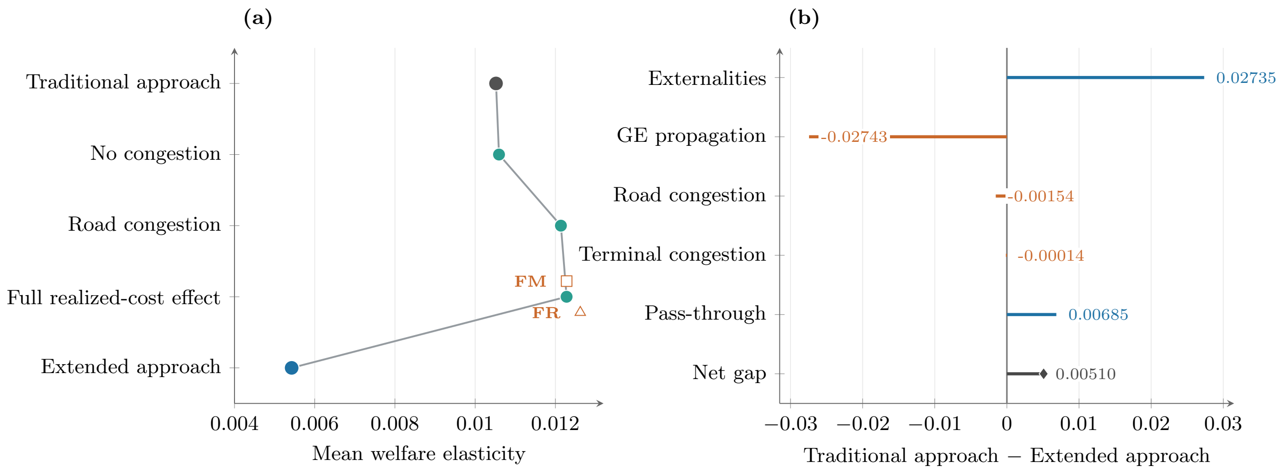

Figure 6: Closure ladder and exact welfare decomposition. Panel (a) moves from the Traditional approach through the no-congestion, road-congestion, and full realized-cost cases to the Extended-approach primitive-cost effect. The hollow FM and FR markers hold modes and routes fixed. Panel (b) reports the signed contributions to Traditional-approach welfare minus primitive-cost welfare; positive values widen the gap and negative values offset it.

Panel (a) of Figure 6 shows which additions change the mean welfare elasticity. Panel (b) accounts for the gap between the Traditional approach and the primitive-cost mean. The common externality term is 0.027352, but general-equilibrium propagation contributes -0.027426 and nearly cancels it. Road and terminal congestion contribute -0.001543 and -0.000142. Primitive-cost pass-through contributes 0.006855 in the opposite direction. The components sum exactly to the mean difference between the Traditional and Extended approaches.

The cancellation reflects opposing economic effects rather than a failed identity: the reported residuals are below the stated tolerance. Spatial externalities generate a common scale, whereas market-access propagation depends on each link's position in the network, so the two contributions can have opposite signs. The road and terminal terms measure changes in congestion exposure. Pass-through measures a separate margin: induced traffic partly offsets the primitive cost reduction before it reaches aggregate edge costs. Reporting signed levels displays these offsets without dividing by a net gap that is small relative to its components.

6.6 Adapting the example

First reproduce the committed grid before replacing its inputs. Section 1 then gives the complete workflow for creating an external data project, updating the package without overwriting that project, validating the new baseline, and recording a new result vintage. The generic loader requires model-ready flows and does not automatically balance, symmetrize, or rescale them.

Other examples exercise specialized parsers and geometries and are catalogued in the repository's examples index. Subsection 5.3 defines the distinction between numerical verification and external validation that also governs those examples.

7 Applications and extensions

Applications differ in their spatial margins, available modes, and empirical flow construction, even though they use the same transport operators. The U.S. freight exercise shows the complete economic-geography application. The urban material that follows changes the spatial equilibrium while retaining the recursive routing, modal, and congestion maps.

7.1 U.S. freight

The freight application maps domestic and international value flows onto a multimodal U.S. network and evaluates bidirectional road improvements. We first define the exercise and its flow construction before reporting the specification, results, decomposition, and provenance.

7.1.1 Purpose

The paper applies the statistic to the U.S. freight network (Allen et al., 2026). The archived paper-replication result has 234 nodes: 228 domestic locations and six foreign regions. The policy set contains 704 directed Interstate road arcs paired into 352 bidirectional physical links. Each experiment lowers the primitive road cost by one percent in both directions of one physical link.

The exercise is a sufficient-statistic application. It does not estimate a new 234-location structural model from raw microdata. The adapter combines model-ready domestic trade and modal flow matrices, location shares, the road network, and directional 2017 Census port trade. It then checks and constructs the inputs required by the derivative.

7.1.2 International flows

The headline vintage preserves separate import and export directions from public 2017 Census containerized vessel trade (U.S. Census Bureau, 2026a). Port values are mapped to domestic nodes and six foreign regions. The adapter first normalizes imports and exports separately. It then averages the normalized U.S.-port margins across the two directions and does the same for the foreign-region margins. Iterative proportional fitting projects each directional matrix onto these common margins while preserving its observed zero pattern. The routine requires a maximum margin error below 10^{-12}; it fails if the observed support cannot sustain the common margins or if the iteration does not converge. Finally, the adapter assigns one-half of the normalized foreign-water modal mass to each direction. This balancing step is part of the model adapter. It does not assert that raw port-level imports and exports are equal.

The raw monthly API cache is not committed. The repository instead provides a credential-safe downloader, port and foreign-region crosswalks, and the aggregate derived overlay. It also records source metadata and hashes, along with diagnostics for coverage, iteration counts, margin errors, total variation, and node-level residuals. The workflow reads the API key from the environment. Its tests and secret scan verify that the tracked cache metadata, URLs, manifests, and outputs contain no key.

7.1.3 Headline specification

The archived configuration uses:

$$\alpha = 0.10, \quad \beta = -0.30, \quad \sigma = 9, \quad \eta = 1.099, \quad \gamma_{\text{road}} = 0.092.$$

Subsection 5.1 documents the corresponding calibration sources: Allen and Arkolakis (2014, 2020) and Fuchs and Wong (2026). It uses `ChoiceLogsum`, all observed transport modes in equilibrium, edge-local road congestion, and no headline terminal congestion. The configuration field retains the historical name `lambda_road`; γ_{road} is the theoretical congestion elasticity used in this guide. These parameters imply $e = -0.5$ and $\rho = -1.6$. The common factor ρ scales all links equally. Traffic, pass-through, and the two market-access multipliers generate the cross-link ranking. The sign of ρ alone does not determine the welfare sign because the multiplier sum is also signed; the archived baseline produces positive benefits.

The policy is a primitive road-cost improvement. The Traditional approach uses the two-direction road traffic share. The full realized-cost effect changes the cost users face after congestion. The Extended approach reports the full primitive-cost effect and includes route, mode, congestion, and spatial-equilibrium responses.

7.1.4 Archived paper-replication results

The current package reproduces the archived numerical reference:

Table 2: Archived physical-link result summary

Measure	Value
Physical links	352
Mean Traditional-approach elasticity	0.0003872
Mean full realized-cost effect	0.0005134
Mean full primitive-cost effect	0.0002231
Median primitive-cost effect	0.0001807
Maximum primitive-cost effect	0.0008719
Pearson correlation with Traditional approach	0.947
Spearman rank correlation	0.959

Implementation. These values pass the package's accounting, algebraic, finite-difference, and reproducibility checks. Those checks verify the baseline construction and the derivative calculation. They do not validate the calibrated congestion, modal, or spatial responses against untargeted empirical outcomes.

The high correlations show that traffic tracks the Extended approach closely across links. They do not imply identical targeting. Only 4 links overlap between the two top-ten lists in the archived result vintage. The paper reports names only where endpoints can be assigned by point-in-polygon matching to public 2018 Census CBSA geography.

7.1.5 Signed decomposition

To explain why the mean primitive-cost derivative differs from traffic, apply the signed Hulten-gap identity in Subsection 4.3 and average each level contribution across physical links:

Externality scale	0.0010068
Equilibrium propagation	−0.0011017
Road congestion	−0.0000313
Terminal congestion	0.0000000
Primitive-cost pass-through	0.0002903.

The externality and propagation terms have opposite signs, and each is large in absolute value relative to their sum. They therefore nearly offset. The full realized-cost mean exceeds the Traditional approach, whereas primitive-cost pass-through lowers the headline primitive result. The decomposition reports signed welfare levels because dividing these terms by the small net gap would produce unstable shares.

7.1.6 Interpretation

The result ranks local model-implied benefits, not project net present values. A high-ranked link may be expensive to build, face environmental constraints, or conflict with distributional objectives. Conversely, a low-ranked capacity improvement may be justified by safety or resilience benefits outside the model. The appropriate use is to add the spatial-equilibrium benefit measure to a broader appraisal, not to suppress the rest of the appraisal.

7.1.7 Result provenance

The guide asset script reads a public-safe aggregate replication record and checks it against the declared configuration, input, source-result, and output hashes in `replication/rsue/expected_summary.toml`. With restricted results present, it also verifies the full directed and physical output files. The archived vintage requires the stated network counts, the numerical gates, a hash-bound finite-difference report, and deterministic outputs. Appendix F records the reproducibility boundary; the repository audit contains the developer-level kernel lineage.

7.1.8 Reproducing with external inputs

Set the external data root and run the locked environment:

```
export RSUE_DATA_ROOT=/path/to/hash-verified/rsue-inputs
julia --project=replication/rsue/environment \
  -e 'using Pkg; Pkg.instantiate()'
julia --project=replication/rsue/environment \
  replication/rsue/verify_choice_logsum_fd.jl \
  replication/rsue/rsue_paper_choice_edge_census_2017.toml \
  replication/rsue/verification/choice_logsum_rsue_fd.toml
julia --project=replication/rsue/environment \
  replication/rsue/build_paper_artifacts.jl
```

The manifest rejects files whose hashes differ from the declared data vintage. There is no fallback to the legacy port matrix.

The freight application keeps the paper's economic-geography model. The urban application instead changes the spatial residuals and welfare projection while retaining the same route, modal, and congestion operators. The next subsection describes its implementation and data requirements.

7.2 Urban commuting and public transit

The freight application uses goods-market clearing and mobile labor to map transport costs into welfare. Urban applications instead use residence and workplace choices. Because both settings share the transport block, the extension changes the spatial residuals and welfare projection while leaving route and mode aggregation intact.

7.2.1 Urban implementation and verification

The urban application changes the spatial model, not the transport technology. In the freight model, locations produce and consume traded goods. In the urban model, workers choose residences and workplaces and travel between them. These choices produce different equilibrium residuals and a different welfare row. For any given residence–workplace flow, however, commuters face the same recursive routes, modal logsum, and congestion feedback as shipments in the freight model.

Section 3 defines the residence–workplace model, and Subsection 4.2 shows how its Jacobian and welfare row enter the shared transport derivative. The implementation stores those model-specific quantities in `UrbanEngine`. The route, modal, congestion, and primitive pass-through maps remain in the common transport modules.

The one-mode regression oracle independently codes the Allen–Arkolakis urban derivative (Allen and Arkolakis, 2022). The shared engine must reproduce its state response and welfare derivative below 10^{-10}, including the $1 + \theta_U \gamma$ congestion denominator. To see the role of this denominator, consider one road mode with edge-local congestion elasticity γ. A lower primitive edge cost attracts additional traffic, which raises the realized cost. Route demand responds with curvature θ_U, so the congestion feedback attenuates the primitive shock by the factor $1/(1 + \theta_U \gamma)$. Reproducing this factor verifies that the shared transport operator has the same local feedback as the original one-mode urban system.

The multimodal urban example adds road and transit to the same model and compares the analytical IFT with an independently solved nonlinear system under `NC`, `NT`, `F`, `FM`, and `FR`. The tests allow residence and workplace margins to differ and include both road and transit modes under edge and endpoint-terminal congestion. They also verify the one-mode, unique-route, no-congestion, and zero-terminal limits, as well as invariance to node and mode permutations.

Run the self-contained example with:

```
julia --project=. bin/tnw.jl validate \
  examples/urban_multimodal/config.toml
```

```
julia --project=. bin/tnw.jl decompose \
  examples/urban_multimodal/config.toml
```

7.2.2 Endpoint-terminal congestion

Edge congestion assigns traffic feedback to a particular edge-mode. Terminal congestion aggregates affected flows using a terminal and applies the cost response to incident edge-mode pairs. The terminal identifiers are explicit input fields; geographic co-location is not enough.

In the freight extension, rail traffic loads the terminals at both endpoints with elasticity 0.096, estimated by Fuchs and Wong (2026). The main paper result excludes terminal congestion. Sensitivity to terminal congestion is therefore an extension diagnostic, not a headline parameter sensitivity.

7.2.3 Seattle multimodal candidate

Practical caution. The Seattle module is a candidate empirical application. Its multimodal urban derivative and accounting checks are implemented, but route ridership, generalized modal costs, Seattle-specific substitution, and congestion parameters require additional evidence before the results can be used as an empirical appraisal.

The Seattle workflow combines the published 217-location Allen–Arkolakis urban geography (Allen and Arkolakis, 2022) with 2017 LODES residence–workplace commuting (U.S. Census Bureau, 2026b) and 2017 ACS residential commute-mode shares (U.S. Census Bureau, 2017). A hash-pinned June 2017 King County GTFS feed supplies the transit network and schedules (King County Metro, 2017a). The Metro system evaluation enters only as an external comparison (King County Metro, 2017b).

The baseline uses $\eta = 1.099$, transferred from the freight application in Fuchs and Wong (2026). It is not estimated from Seattle data. Baseline modal shifters reproduce the constructed shares; η controls the local response after an improvement. The workflow therefore reports sensitivity at 0.75, 0.90, 1.099, 1.25, and 1.40.

GTFS describes stops, trips, shapes, and schedules. It does not provide passenger counts. ACS supplies residential mode shares, not route-level ridership. The builder routes a common commuter population to obtain a model-consistent road and transit baseline, records fallback and coverage, and uses the Metro table only as an external comparison.

7.2.4 Transit policies

The workflow evaluates bus, Link light rail and streetcar, and ferry policies separately and together. A link policy lowers primitive generalized cost on one directed transit edge-mode. Reciprocal pairs are aggregated only when both directions exist.

A named route corridor changes several edge-mode costs at once. Let b denote its set of edge-mode pairs, $w_{e,m}$ the relative cost-reduction weight assigned to pair (e, m), and $E_{e,m}$ its marginal welfare elasticity. Linearity of the first-order derivative gives the elasticity of the bundled reduction:

$$E_b = \sum_{(e,m)\in b} w_{e,m} E_{e,m}.$$

Shared segments change the aggregate mode cost on that segment. The result is not interpreted as a shock exclusive to one operator's vehicles.

Maps distinguish road, bus, Link light rail, Seattle Streetcar, and ferry by color. Line width represents the primitive-cost welfare gain, not ridership. Subway or light-rail vision maps that include proposed future lines are not a baseline description of the 2017 network.

7.2.5 Remaining empirical work

The package computes and verifies the multimodal urban derivative, but a publication-quality Seattle transit application still requires observed passenger counts by route or segment and mode-specific measures of generalized cost and service quality. It also needs a Seattle-specific estimate or external calibration of η, evidence on crowding or terminal-congestion elasticities, and validation against commuting and ridership outcomes that were not used to construct the baseline. These requirements concern measurement and calibration, not the IFT method. Until they are met, the Seattle module tests the implemented residence–workplace derivative and multimodal transport operators but does not provide an empirical transit appraisal. These empirical boundaries also govern how results from the Seattle module should be described. Section 8 explains how they enter interpretation and reporting.

8 Interpretation and reporting

A transport improvement first changes the generalized cost of using a link. Traffic on that link measures the direct saving. At an efficient interior baseline, the envelope theorem makes this traffic-based statistic the complete first-order welfare effect: subsequent changes in routes, modes, prices, and locations have no additional first-order value. This result explains both the appeal of social-savings calculations and the conditions under which they are sufficient.

The Extended approach is needed when the baseline contains congestion or spatial externalities. Route and modal responses then move activity across wedges, and the primitive policy shock need not pass one for one into the realized transport cost. Proposition 2 organizes these adjustments as traffic times pass-through and market-access multipliers. The adjoint calculation implements the same derivative for the declared set of candidate links without resolving a separate counterfactual equilibrium for each one. The model decomposition then asks which modeled effect accounts for the difference from the Traditional approach.

The economic-geography and urban-commuting applications share this transport calculation. They differ in the spatial equilibrium that maps transport costs into production and consumption or into residence and workplace choices. Consequently, a practitioner must choose the spatial model before interpreting the welfare result. The same network and traffic data can imply different welfare derivatives when the underlying economic adjustment margin changes.

Three checks support different claims. Accounting checks establish that the input flows, routes, modes, and policy units form a coherent baseline. Derivative and finite-difference checks establish that the code computes the stated model. Comparisons with untargeted flows, costs, or outcomes assess whether that model fits a particular application. Passing one class of checks does not substitute for passing the others.

The output is a local model-implied benefit from a specified transport-cost change. It is not a complete project appraisal. Construction and financing costs, safety, local pollution, displacement, and distributional incidence must be added separately when they matter for the decision. A reproducible run therefore reports the policy unit, economic model, parameters, input transformations, diagnostics, and artifact hashes alongside the welfare elasticity. Reported in this form, the statistic extends traffic-based accounting by measuring the spatial-equilibrium correction implied by the stated model and data. The reference appendices provide the notation-to-code crosswalk, algorithms, schemas, numerical tests, example results, and provenance needed to audit that calculation.

Part IV Reference appendices

The main guide emphasizes the sequence from policy definition to an interpretable result. The appendices collect the notation, algorithms, schemas, tests, additional examples, and provenance needed to reproduce or extend a particular stage without interrupting that sequence.

A Notation and equation-to-code crosswalk

Sections 3 and 4 present the two spatial models and their common transport calculation in economic notation. This appendix records the corresponding code quantities and fixes the orientation and sign conventions used by the algorithms in Appendix B.

A.1 Core notation

Symbol or term	Economic meaning	Code representation
i, j, k, l	Location and directed-edge endpoints	Integer node indices plus string node IDs
m	Transport mode	`Symbol` and mode index
$\bar{\kappa}_{ij,m}$	Primitive edge-mode iceberg cost	Policy shock selector
$\kappa_{ij,m}$	Realized edge-mode cost after congestion	Congestion state map
κ_{ij}	Aggregate directed-edge modal logsum	`Sagg` and modal response
τ_{ij}	Bilateral route cost index	Route resolvent and soft-route operators
$\Xi_{ij,m}$	Edge–mode traffic: world-income-normalized value flow in economic geography or normalized commuter flow in the urban model	Mode-flow sparse matrix
Ξ_{ij}	Aggregate edge traffic	`data.Xi`
$s_{ij,m}$	Modal traffic and cost-elasticity share	`data.s_edges` and `basis.Sagg`
$\mathcal{P}_i, \mathcal{Q}_i$	Inward and outward market-access stocks	Reconstructed exposure quantities
$\mathcal{T}_i^G$	Trade exposure before local absorption	`node_visit_stock`; technical legacy name
$\mathcal{T}_i^U$	Urban commuting exposure satisfying $l_i^F + O_i = l_i^R + I_i$	`node_visit_stock` under urban loading
l_i^R, l_i^F	Urban residence and workplace margins	`data.residence`, `data.workplace`
$z_G \in \mathbb{R}^{2N}$	Economic-geography state, ordered as the column vector of outward and inward transformed log variables	`AdjointRSUE` state vector
$z_U \in \mathbb{R}^{2N+1}$	Urban state, ordered as residence logs, workplace logs, and the welfare normalizer	`UrbanEngine` state vector
ϖ_U	Urban equilibrium normalizer, with $d \log W = -d \log \varpi_U / \theta_U$	Final urban state component; historical code name `chi_U`
μ_{ij}	Outward continuation share	`data.mu`
λ_{ij}^{in}	Inward continuation share	`data.lam`
s_i^x, s_i^y	Local outward and inward absorption shares	`data.sx`, `data.sy`

Symbol or term	Economic meaning	Code representation
α, β	Model-specific spatial response parameters: productivity and amenity elasticities in economic geography; workplace and residence response parameters in the urban model	`StructuralParameters`
σ	Trade elasticity	`parameters.sigma`
θ_U	Urban commuting heterogeneity and route curvature	`model.theta`, stored internally through $\sigma - 1$
η	Modal substitution parameter	`ChoiceLogsum.eta`
ω	Route curvature	$\sigma - 1$ or urban $\theta_U =$ `model.theta`
θ_{ijm}	Primitive edge-mode log-cost policy	Policy selector and `primitive_forcing`
$\gamma_{e,m}$	Edge–mode congestion elasticity	Typed congestion specification; legacy configuration fields use `lambda` names
e	Transformation regularity scalar	`AdjointRSUE.Coef.e`
ρ	Externality and welfare scaling coefficient	`AdjointRSUE.Coef.rho`
$J_{c,0}, J_{c,S}$	Model-specific no-congestion and full spatial Jacobians	`closures.J0`, `closures.F`, etc.
q_c	Welfare gradient in the selected spatial state	`welfare_gradient`
ψ_c	Adjoint solution to $J_c^\top \psi_c = q_c$	Transposed linear solve in `operator_gain`
b^r, b^θ	Realized and primitive shock forcing	`realized_forcing`, `primitive_forcing`
ζ_{ijm}, χ_{ijm}	Aggregate-cost pass-through and its modal-share normalization	`primitive_forcing`
χ_e^{eff}	Welfare-effective primitive-to-realized ratio	`chi_effective` where defined

A.2 Main equations

Theory quantity	Construction	Primary implementation
Modal logsum	Negative-power aggregation and share derivative	`modal_power`, `pair_response_matrices`
Recursive routing and curvature	$T = (I - K)^{-1}$, with $\omega = \sigma - 1$ in the paper	`reconstruct_route_kernel`, `route_curvature`
Bilateral reconstruction	$\operatorname{diag}(o) T \operatorname{diag}(a)$	`soft_route_operators`
Edge–mode traffic	$\Xi_{e,m} = s_{e,m} \Xi_e$	`mode_flows`, `Sagg`
Traffic factorization	Continuation share times exposure	`shares_from_traffic`, route diagnostics
Transformed inverse map	(x, y) to economic variables	`inverse_state_map`
Economic-geography spatial Jacobian	Sparse four-block system plus aggregate feedback	`AdjointRSUE.assemble_J`
Urban spatial Jacobian	Residence–workplace system and two margin normalizations	`UrbanCommutingIFT.jacobian`, `UrbanEngine`
Congestion elimination	Schur complement	`build_transport_closure`
Adjoint derivative	$q_c^\top J_{c,S}^{-1} b_c$	`operator_gain`

Theory quantity	Construction	Primary implementation
Realized and primitive forcing	b^r versus b^θ	`realized_forcing`, `primitive_forcing`
Primitive pass-through	ζ from the transport response and $\chi = \zeta/s$	`primitive_forcing`, `effective_ratio`
Paper Proposition 1	Efficient derivative equals value traffic	Hulten tests and `decomposition_rows`
Paper Proposition 2	Traffic times ρ, χ, and access multipliers	Adjoint projection; scalar helper for oracle tests
Urban welfare derivative	$q_U^\top J_{U,S}^{-1} b_{U,p}$	`UrbanEngine`, `operator_gain`
Closure ladder	`NC`, `NT`, `F`, `FM`, `FR`	`build_closures`
Inverse-gap identity	Difference between closure derivatives	`inverse_gap`
Mixed decomposition	General allocation, scarcity, and equilibrium formula; the current common-baseline specification leaves only scarcity nonzero	`mixed_channels`
Physical aggregation	Sum reciprocal policies before normalization	`aggregate_physical`

A.3 Orientation conventions

An edge tuple (i, j) means origin i, destination j. In the row-oriented continuation matrix μ, the element μ_{ij} is the share leaving i on edge $i \to j$. The inward share is stored using the same edge orientation, with destination exposure in the denominator. Code that differentiates the second market-access recursion uses the transpose where required.

The policy selector indexes positive-flow edge-mode pairs rather than all possible matrix cells. The directed result retains both integer indices and source IDs. Physical aggregation relies on endpoint IDs and the declared `physical_link_id`, not adjacency-matrix symmetry alone.

A.4 Signs

Costs enter as log increases, while reported benefits correspond to cost reductions:

$$E_p = -\frac{d \log W}{d \log \bar{\kappa}_p}.$$

The Hulten traffic share is positive. The primitive and realized elasticities are stored as positive benefits under this convention. Decomposition wedges and decomposition components remain signed and can be negative.

The coefficient ρ can be negative. Do not replace signed products by absolute values. A negative ρ and negative access multiplier sum can produce a positive benefit.

A.5 Legacy aliases

The mathematical audit and early code use "node-visit stock" or "occupation stock" for the recursive exposure quantity. These are technical descriptions of the resolvent. The paper and guide use "trade exposure" for $\mathcal{T}_i^G$ and "commuting exposure" for $\mathcal{T}_i^U$. The code function name is preserved for compatibility.

The positive-power modal type `ComponentCES` is **legacy-only**. Its congestion helpers are not part of the active `ChoiceLogsum` implementation. The distinction changes the economics of modal substitution;

it is not a cosmetic sign convention.

B Computational algorithm

Subsection 5.2 describes the package architecture. This appendix follows one policy derivative through validation, route reconstruction, closure construction, the adjoint solve, and the decomposition. The pseudocode omits storage and dispatch details but preserves the mathematical order of operations.

B.1 Validation algorithm

Validation must establish a common baseline before any Jacobian is assembled. The algorithm first checks schemas and accounting, then reconstructs routes, and only then tests the full welfare closure:

```
function validate_project(config)
    project = load_project(config)
    data = load_network(project)

    check_schema(data)
    check_mode_and_congestion_keys(project, data)
    check_physical_policy_pairs(project, data)
    check_flow_and_exposure_accounting(data)

    basis = reconstruct_route_basis(project, data;
                                    include_fixed=false)
    require(basis.spectral_radius < 1)
    require(basis.margin_error <= tolerance)
    require(basis.edge_error <= tolerance)

    closure = build_full_welfare_closure(project, data, basis)
    require(condition_number(closure) <= condition_limit)
    return validation_summary(project, data, basis, closure)
end
```

Validation builds only the quantities needed for the welfare calculation. It does not construct the full fixed-route incidence matrix.

B.2 Closure-selected welfare algorithm

For welfare analysis, the route basis supplies the transport derivatives and the selected spatial builder supplies $J_{c,0}$, q_c, and $B_{\kappa,c}$. The Schur complement in (4.5) combines them before the adjoint projection in (4.4):

```
function all_welfare_effects(project)
    data = load_network(project)
    basis = reconstruct_route_basis(project, data;
                                    include_fixed=false)

    J0, q, B = build_spatial_system(project, data)

    transport_F = transport_closure(project, data, basis,
                                    flexible_routes=true,
                                    flexible_modes=true,
                                    edge_congestion=true,
                                    terminal_congestion=true)
    JF = J0 + transport_F.feedback
```

```
    b_primitive = primitive_policy_forcing(B, transport_F, basis)
    adjoint = solve(transpose(JF), q)
    elasticities = transpose(adjoint) * b_primitive

    return directed_and_physical_results(elasticities, data, project)
end
```

For `economic_geography`, `build_spatial_system` constructs the transformed trade system and common-utility welfare row. For `urban_commuting`, it constructs the residence–workplace Jacobian and the welfare row whose last entry is $-1/\theta_U$. The remaining transport and adjoint steps are unchanged.

The implementation uses sparse matrices where possible. The schematic `solve` represents a factorization and backsolve, not an explicit inverse.

B.3 Full decomposition algorithm

Decomposition differs from ordinary analysis because it must construct the fixed-route basis and all common-baseline closures. It then applies the same realized forcing to each closure and reserves the primitive forcing for the full model:

```
function decompose(project)
    data, basis = build_route_basis(project; include_fixed=true)
    J0, q, B = build_spatial_system(project, data)

    NC = closure(J0, no_congestion,
                 flexible_routes, flexible_modes)
    NT = closure(J0, edge_congestion,
                 flexible_routes, flexible_modes)
    F  = closure(J0, full_congestion,
                 flexible_routes, flexible_modes)
    FM = closure(J0, full_congestion,
                 flexible_routes, fixed_modes)
    FR = closure(J0, full_congestion,
                 fixed_routes, flexible_modes)

    b_realized = realized_policy_forcing(B, basis)
    b_primitive = primitive_policy_forcing(B, F, basis)

    E = adjoint_gains((NC, NT, F, FM, FR), q, b_realized)
    Eprimitive = adjoint_gain(F, q, b_primitive)

    check_inverse_gap_identities(E, (NC, NT, F, FM, FR))
    if project.spatial isa EconomicGeography
        channels = analytical_mixed_updates(F, FM, FR)
        check_channel_reconstruction(channels)
    else
        channels = signed_closure_gaps(E)
    end
    check_hulten_accounting(E, Eprimitive, traffic)

    return write_decomposition(E, Eprimitive, channels)
end
```

Both models report exact full-versus-no-congestion, fixed-mode, and fixed-route gaps. The finer allocation, scarcity, and equilibrium-correction split is derived and reported only for economic geography.

B.4 Analytical decomposition of each difference

The economic-geography decomposition separates each closure Jacobian into analytically constructed blocks:

$$J_{G,S} = D_S + u_S v_S^\top + \Delta_{edge,S} + \Delta_{terminal,S}.$$

The sparse block D_S contains local allocation and route shares. The rank-one term imposes aggregate labor feedback and normalization. The final two terms contain edge and endpoint-terminal congestion. These blocks come from the differentiated residual equations rather than a numerical matrix factorization fitted after the fact.

The common-baseline construction leaves D_S, u_S, and v_S unchanged across the five equilibrium closures. For road and terminal comparisons, the code evaluates the inverse difference with a structured parallel-sum formula. Their allocation and aggregate-equilibrium components are exact zeros, so the congestion-scarcity component carries each wedge.

The code evaluates mode and route comparisons through the general mixed update rather than assigning their zero components by convention. Let $C_{A:F}$ be the sparse allocation update from the full closure toward $A \in \{FM, FR\}$, and let $S_{A:F}$ be the corresponding congestion-scarcity update. After applying those two updates, define

$$\widetilde{K}_{A:F} = J_F + S_{A:F} + C_{A:F}.$$

The remaining aggregate-feedback difference has the analytical rank-two form

$$J_A = \widetilde{K}_{A:F} + R_{A:F} Q_{A:F}^\top.$$

Sequential use of the inverse-gap identity and the Woodbury formula then gives

$$d_r = d_r^{alloc} + d_r^{scar} + d_r^{eq}, \qquad r \in \{mode, route\}.$$

The three terms respectively change the sparse allocation core, congestion exposure, and aggregate feedback. The implementation evaluates the general formulas even though $C_{A:F} = 0$ and $R_{A:F} Q_{A:F}^\top = 0$ in the current specification. It rejects a result if the Jacobian reconstruction, rank-two identity, implied zero, or component sum exceeds tolerance.

B.5 Route reconstruction

Let K be the continuation matrix, o the source margin, and a the destination absorption vector. For distinct source and destination margins, the reconstruction used in Subsection 3.5 is:

```
K = continuation_matrix(edge_traffic, exposure)
T = inverse(I - K)
Xod = Diagonal(source_margin) * T * Diagonal(absorption)

require(maximum(abs.(sum(Xod, dims=2) - source_margin)) <= tol)
require(maximum(abs.(sum(Xod, dims=1)' - destination_margin)) <= tol)
require(edge_incidence(Xod, K) matches observed_edge_traffic)
```

The code uses linear solves rather than materializing `inverse(I-K)` when possible. The notation displays the resolvent identity directly.

B.6 Primitive forcing

Let p index active edge-mode pairs and let D inject policy shocks into those pairs. The matrix S_{agg} aggregates pair costs to directed edges, C^{route} maps aggregate edge costs into route responses, and L lifts aggregate edge responses back to edge-mode traffic at baseline modal shares. With modal power $p_m = -\eta$, the direct edge-mode traffic response to all primitive policy shocks is

$$\xi_\theta^{direct} = L(C^{route} S_{\text{agg}} D) + p_m (D - L S_{\text{agg}} D).$$

The second term is the within-edge modal substitution response. Let A aggregate edge-mode traffic into congestion quantities, X_q map congestion costs back into edge-mode traffic, and G map congestion quantities into realized edge-mode costs. Substituting the traffic response into the congestion equation gives

$$g_\theta = (I - A\mathsf{X}_q)^{-1} A \xi_\theta^{direct}.$$

The resulting aggregate directed-edge cost response is

$$h_\theta = S_{\text{agg}}(D + G g_\theta),$$

so the primitive forcing in the spatial residual system is

$$b^\theta = B_\kappa h_\theta.$$

These substitutions are implemented in `primitive_forcing`. When several edges share a terminal congestion state, g_θ can change costs away from the improved edge; primitive pass-through is then nonlocal.

B.7 Complexity

Let N be nodes, E active directed edges, P active edge-mode pairs, and Q congestion states. The welfare path requires route reconstruction, an $O(N)$-dimension spatial factorization, a Q-scale transport solve, and sparse products for all policy pairs. The full route decomposition additionally stores origin-destination edge incidence, which can scale as N^2E.

The validator reports a preflight memory estimate. For large assignment networks, use `analyze` unless the fixed-route comparison is essential and the incidence matrix is computationally feasible.

C Input schemas and configuration

Subsection 5.1 states the economic requirements on inputs and policies. This appendix gives the corresponding TOML fields and CSV columns. The examples declare the model specification directly. They do not encode an undocumented preprocessing step.

C.1 Minimal economic-geography configuration

```
schema_version = 1
name = "my-network"

[input]
adapter = "generic_csv_v1"
nodes = "data/nodes.csv"
edge_modes = "data/edge_modes.csv"

[model]
spatial_specification = "economic_geography"
alpha = 0.10
beta = -0.30
sigma = 9.0
modal_specification = "choice_logsum"
eta = 1.099
```

```
route_curvature = "theorem"

[congestion]
specification = "edge"
[congestion.edge]
road = 0.092

[policy]
mode = "road"
unit = "both"
shock_fraction = 0.01

[output]
directory = "output"

[diagnostics]
tolerance = 1.0e-10
condition_limit = 1.0e12
```

The numerical values in this template use the paper calibration documented in Subsection 5.1; they are not defaults for an unrelated application.

C.2 Spatial specification

`economic_geography` requires α, β, σ, with route curvature $\sigma - 1$. `urban_commuting` instead requires α, β, θ_U, where θ_U is supplied by the TOML field `theta`. Internally, $\sigma = \theta_U + 1$ supplies the shared route curvature. Residence and workplace columns replace labor and income shares. The complete public urban configuration is `examples/urban_multimodal/config.toml`.

C.3 Modal specification

`choice_logsum` is the negative-power modal cost index used by the active paper and urban extension. `component_ces` is a **legacy-only** positive-power component index and should be used only for declared provenance replications. Both require finite $\eta > 0$.

C.4 Congestion specification

None

```
[congestion]
specification = "none"
```

Mode-level edge congestion

```
[congestion]
specification = "edge"
[congestion.edge]
road = 0.092
```

Edge-specific congestion

```
[congestion]
specification = "edge"
```

```
source = "input_column"
column = "congestion_elasticity"
scale = 1.0
```

Each active edge-mode row must contain a finite nonnegative value in the selected column. Mode-level and column-based values cannot be mixed.

Endpoint-terminal congestion

```
[congestion]
specification = "endpoint_terminal"
endpoint_scale = 1.0
[congestion.terminal]
rail = 0.096
```

Each positive-congestion pair must identify both endpoint terminals.

Composite congestion

```
[congestion]
specification = "composite"
endpoint_scale = 1.0
[congestion.edge]
road = 0.092
[congestion.terminal]
transit = 0.04
```

C.5 Policy units

`directed_arc` writes directed results only. `physical_link` requires reciprocal pairs and writes physical-link results only, while `both` writes both result sets.

C.6 Economic-geography node schema

Column	Required	Meaning
`node_id`	Yes	Stable unique identifier
`labor`	Yes	Positive normalized labor share
`income`	Yes	Positive normalized income share
`longitude`	No	Finite map coordinate
`latitude`	No	Finite map coordinate

C.7 Urban node schema

The urban adapter uses the same `node_id`, `longitude`, and `latitude` columns as the economic-geography adapter. It replaces `labor` and `income` with the two margins below.

Column	Required	Meaning
`residents`	Yes	Normalized residence margin l^R
`employment`	Yes	Normalized workplace margin l^F

C.8 Edge-mode schema

Column	Required	Meaning
`edge_id`	Yes	Directed edge identifier, shared by modal rows on the same directed connection
`physical_link_id`	Yes	Identifier used for reciprocal aggregation
`origin`	Yes	Existing node ID
`destination`	Yes	Existing node ID, distinct from origin
`mode`	Yes	Nonempty mode name
`flow`	Yes	Finite positive normalized value flow
`origin_terminal_id`	When used	Required for endpoint-terminal congestion
`destination_terminal_id`	When used	Required for endpoint-terminal congestion
`congestion_elasticity`	When used	Required when selected as the edge-congestion source

C.9 Sensitivity

```
[sensitivity]
alpha = [0.06, 0.08, 0.10, 0.12, 0.14]
beta = [-0.38, -0.34, -0.30, -0.26, -0.22]
eta = [0.75, 0.90, 1.099, 1.25, 1.40]
lambda_road = [0.00, 0.04, 0.092, 0.14]
```

Column-based congestion uses `edge_congestion_scale` rather than a mode-specific lambda path.

D Numerical verification

Subsection 5.3 explains how to interpret accounting, algebraic, and empirical checks. This appendix records the tests and numerical gates used to accept the first two classes. Passing them establishes fidelity to the specified model; it does not establish external validity.

D.1 Acceptance tolerances

Table 3: Default verification gates

Check	Tolerance
Efficient Hulten collapse	10^{-8}
Algebraic closure and decomposition identities	10^{-10}
Route margins and edge reconstruction	10^{-10}
Choice-logsum nonlinear finite differences	10^{-6} absolute
Nonlinear equilibrium residual	10^{-10}
Maximum accepted condition number	10^{12}

D.2 Modal tests

Tests differentiate the nonlinear `ChoiceLogsum` independently and compare it with the analytical modal matrix. A cost reduction must shift traffic toward the improved mode. Mode permutation must leave aggregate results unchanged after labels are restored. `ComponentCES` is also kept under regression test, but only as a legacy specification.

One-mode cases must satisfy $\texttt{FM} = \texttt{F}$. A zero-flow mode is not an interior active pair and cannot be a policy target.

D.3 Route tests

Route tests check contraction, absorption, bilateral margins, and edge traffic. Flexible-route derivatives are compared with central differences of the resolvent. Fixed-route incidence is reconstructed independently.

In a unique-route network, $\texttt{FR} = \texttt{F}$. Parallel-route fixtures must produce a verified nonzero route wedge. Node permutations must preserve the economic result.

D.4 Congestion tests

Edge and terminal congestion are tested separately and jointly. Zero congestion implies:

$$\texttt{NC} = \texttt{NT} = \texttt{F}, \qquad E^{\theta} = E^{r}.$$

Zero terminal congestion implies $\texttt{NT} = \texttt{F}$. Unknown congestion keys, positive values for inactive modes, and missing terminal identifiers must fail during validation.

The grid example checks edge and endpoint-terminal congestion separately. Its zero-terminal version must satisfy $\texttt{NT} = \texttt{F}$. Package examples with edge-specific BPR elasticities retain a separate formula and finite-difference regression test, but they are outside the guide's worked example.

D.5 Spatial and welfare tests

The Python oracle provides 55 checks of the corrected transformed economic-geography system. Julia fixtures compare Jacobian blocks, the welfare row, the Hulten limit, and the legacy result interfaces.

The active choice-logsum path is checked independently using nonlinear equilibria. The freight verification chooses a deterministic traffic-stratified set of arcs, solves positive and negative primitive shocks, and binds the report to configuration, input, and source hashes.

The guide example adds central finite differences for one road improvement and one transit improvement. It also checks its efficient Hulten limit, one-mode limit, zero-terminal limit, route contraction, and full decomposition identities. Unique-route equality is covered by a separate deterministic fixture because a grid intentionally has alternative routes.

D.6 Urban tests

The one-mode urban oracle checks the Allen–Arkolakis state response and welfare derivative. The shared multimodal engine is then compared with `UrbanNonlinear.jl`, which solves the nonlinear equilibrium independently.

Tests cover `NC`, `NT`, `F`, `FM`, and `FR`; distinct residence and workplace margins; road and transit; terminal congestion; and mode and node permutation.

D.7 Output tests

Output tests verify the behavior of directed, physical-link, and combined policy units, including reciprocal aggregation without dropped or duplicated directions. They also check strict JSON treatment of missing values and key collisions, shock-fraction scaling of generated TeX macros, deterministic output hashes, and full Census CBSA labels, including hyphenated names.

D.8 Clean-clone acceptance

A release candidate must pass:

```
julia --project=. -e 'using Pkg; Pkg.instantiate()'
julia --project=. -e 'using Pkg; Pkg.test()'
julia --project=. bin/tnw.jl validate \
  examples/grid_multimodal/config.toml
julia --project=. bin/tnw.jl decompose \
  examples/grid_multimodal/config.toml
julia --project=docs -e \
  'using Pkg; Pkg.develop(path="."); Pkg.instantiate()'
julia --project=docs docs/make.jl
make practitioner-guide-assets
make practitioner-guide-check
python3 scripts/verify_provenance.py
python3 scripts/scan_secrets.py
git diff --check
```

The Julia calculations and PDF build require no restricted inputs, Overleaf, Dropbox, or credentials. Asset regeneration uses the pinned environment in `plots/requirements.txt`; ordinary analysis and PDF compilation do not.

E Examples and application adapters

The worked grid in Section 6 is the shortest route through the entire calculation, but it is not the package's only example. This appendix documents the repository's ten examples, including the external-data builders whose raw inputs remain outside Git. The examples serve different purposes. The toy, Braess, and multimodal urban networks isolate specific parts of the calculation. Sioux Falls adapts a familiar traffic network, while Seattle and Westeros apply the two spatial models to larger geographies. A numerical result appears only when the corresponding input passes the package's accounting and numerical checks.

E.1 Coverage and status

Example	Spatial model	Construction	Status and role
Three-node toy	Economic geography	Hand-written labor, income, road, rail, and terminal records	Self-contained API and sensitivity smoke test
Multimodal grid	Economic geography	Programmatically generated 5×5 road grid with a central transit spine	Self-contained worked example in Section 6
Braess-style network	Economic geography	Hand-written four-node, two-route network with a connector	Self-contained route-closure test
Cow network	Economic geography	Original 30-node three-dimensional geography; optional external PLY mesh	Self-contained ranking example; large mesh is on demand
Sioux Falls	Economic geography	Pinned TNTP network, trips, assigned flows, and local BPR elasticities	Verified external-data benchmark adaptation
Three-location urban toy	Urban commuting	Hand-written residence, workplace, OD, and road-flow records	Self-contained one-mode regression oracle
Multimodal urban toy	Urban commuting	Four locations, road and transit, alternative routes, and terminal congestion	Self-contained shared-transport test
Seattle road diagnostic	Urban commuting	Published LODES and HPMS/HERE inputs from the Allen–Arkolakis replication	Rejected by the accounting gate; no welfare result is reported
Seattle multimodal	Urban commuting	LODES OD, ACS mode shares, and hash-pinned 2017 GTFS service	Verified candidate application, not an externally calibrated appraisal
Westeros	Economic geography	Public ArcGIS geography with synthetic activity, trade, and traffic	Verified illustrative calculation; not an empirical claim

The maps, scatters, and ranked tables below follow one design rule. Map color records the Extended-approach welfare elasticity, line width rises with its absolute value, and node area records economic activity. Each scatter uses the same scale on both axes, so the dashed line marks equality between the Traditional and Extended approaches. The figures generated from external data are frozen to input and output hashes. Rebuilding them requires the source data, but compiling this guide does not.

E.2 Self-contained economic-geography examples

These examples keep their model inputs in the repository and differ in network structure and the economic response they are designed to expose. The smallest case checks the API; the remaining cases add route, modal, and congestion comparisons.

E.2.1 Three-node API toy

The smallest generic example has three locations and three reciprocal physical links. Road is available on all three links, while rail is available on two links and loads an endpoint terminal at each end. The input files record positive labor and income at each location and model-ready value flows by edge and mode. The configuration combines edge-local road congestion with rail-terminal congestion and evaluates a one-percent bidirectional road improvement. Because the files are short enough to inspect directly, this example is the package's first check of parsing, normalization, modal aggregation, terminal identifiers, policy-unit handling, decomposition, and the seven sensitivity paths.

```
julia --project=. bin/tnw.jl validate examples/toy/config.toml
julia --project=. bin/tnw.jl decompose examples/toy/config.toml
```

The example is original to this package. Its coordinates are only plotting metadata, and its flows are economic value shares rather than measurements from the three cities named by those coordinates.

E.2.2 Multimodal grid

The grid used throughout the guide is generated by `examples/grid_multimodal/build_inputs.jl`. Reciprocal road links join adjacent locations on a 5×5 lattice. Transit is available on the four links along the middle row, and explicit terminal identifiers load the terminal-congestion extension. Symmetric directional flows and positive location activity are chosen together so that the recursive-flow identities hold without balancing inside the estimator. Section 6 reports the network map, Traditional–Extended comparison, ranked links, and exact decomposition. The example and all of its displayed assets are original and fully distributable.

E.2.3 Braess-style routing network

The Braess-style example uses the four-node topology associated with the classic traffic-planning paradox (Braess, 1968). Two paths connect the outer locations, and a cross-link joins the intermediate nodes. The package does not solve a Wardrop assignment, so the example does not claim to reproduce Braess's paradox. Instead, reciprocal model-ready value flows are placed on the five physical links. The full flexible-route response can then be compared with `FR`, which holds baseline physical-edge use fixed. This construction isolates the route term while retaining the economic-geography equilibrium.

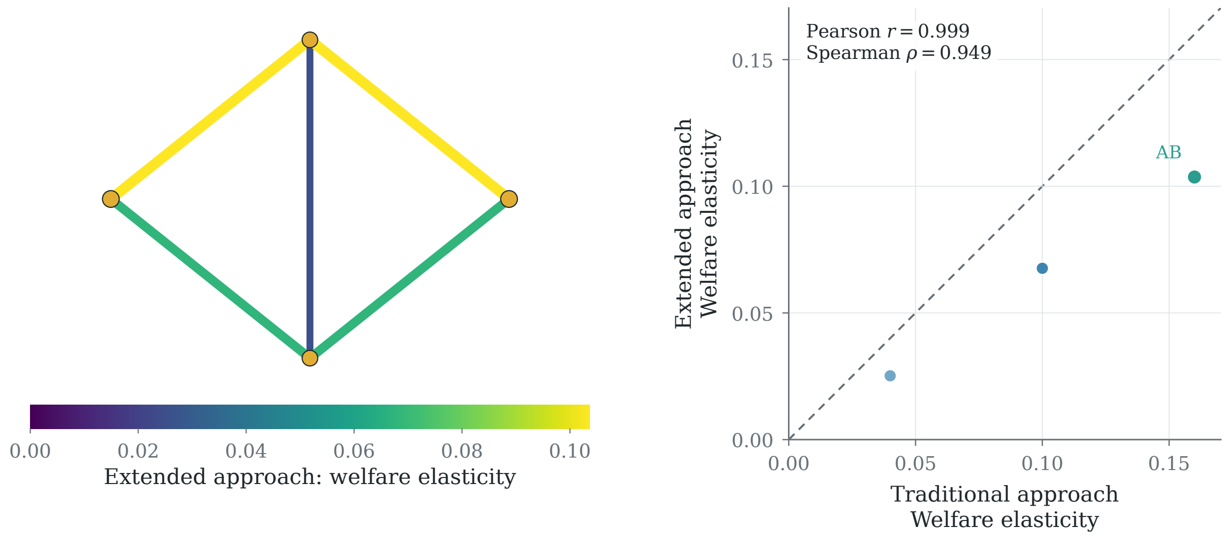

Figure 7: Braess-style routing example. The left panel maps the Extended-approach elasticity. The right panel compares the Traditional and Extended approaches on common axes.

Table 4: Physical links in the Braess-style example

Link	Traditional	Extended	Traditional rank	Extended rank
AB	0.160000	0.103643	1	1
BD	0.160000	0.103643	2	2
AC	0.100000	0.067615	3	3
CD	0.100000	0.067615	4	4
BC	0.040000	0.025182	5	5

The route comparison is nonzero because the connector and the two outer paths give traffic more than one way to reach the destination. All five Extended-approach points lie below the equality line in this calibration. That ordering is a result of the chosen externalities, congestion, and equilibrium response; it is not imposed by the routing operator.

E.2.4 Cow network and full mesh

The 30-node cow is an original synthetic geography. Nodes and links form a three-dimensional body, head, legs, and two alternative neck routes. Positive activity is concentrated in the body and head. Model-ready balanced flows then give the neck a clear economic role: a link can matter because it carries traffic, because it changes access between the two high-activity regions, or because it changes how traffic uses the alternative route. The elevation coordinate affects only the visualization.

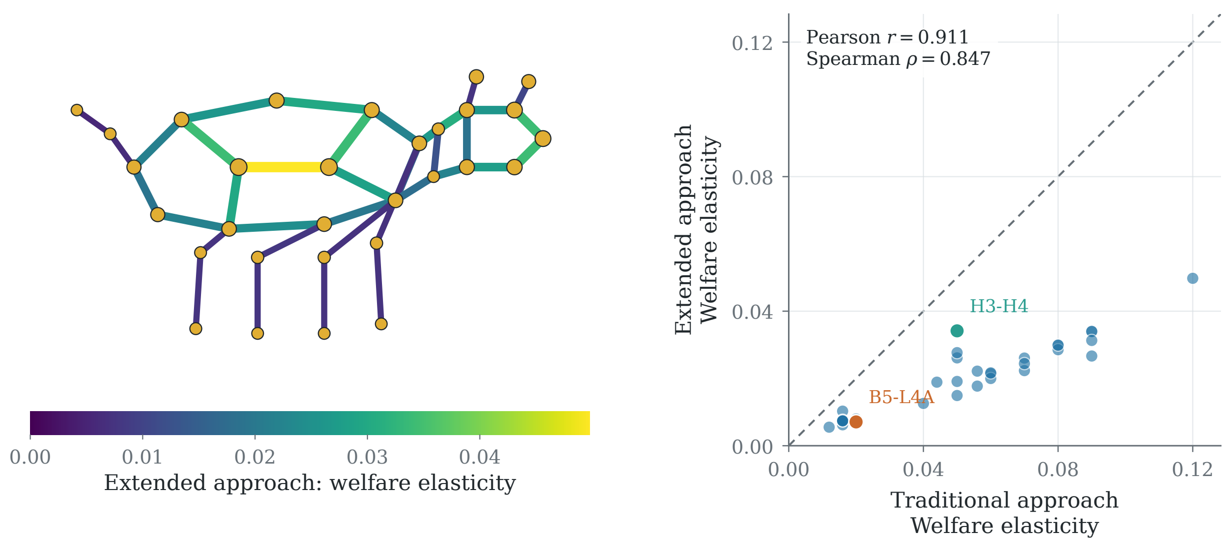

Figure 8: Cow-shaped economic geography. Both panels use the common example design; the three-dimensional rendering supplied by the package uses the same welfare scale.

Table 5: Highest-ranked physical links in the cow-shaped network

Link	Traditional	Extended	Traditional rank	Extended rank
C1-C2	0.120000	0.049769	1	1
H3-H4	0.050000	0.034154	21	2
H2-H3	0.050000	0.034081	20	3
B2-C1	0.090000	0.034050	2	4
B4-C2	0.090000	0.033947	3	5
H1-N1	0.090000	0.031296	5	6
B8-C1	0.080000	0.029924	8	7
B3-B4	0.080000	0.029821	6	8

Only 3 of the Traditional top 5 links remain in the Extended top 5, and the largest rank movement is 19 positions. The closure results contain nonzero congestion and route terms, which are two channels through which the Extended ranking can depart from the traffic ranking. The calculation does not assign each individual rank change to a single channel.

The optional high-resolution version begins with John Burkardt's 2,903-vertex PLY cow (Burkardt, 2026). The builder assigns positive labor and income to all 2,903 mesh vertices, turns each mesh edge into a reciprocal road link, and uses symmetric matrix scaling to construct balanced value flows. The resulting network has 8,706 physical links and 17,412 directed arcs. Its sparse edge-local calculation evaluates the same primitive-cost derivative, but it does not build the dense fixed-route and fixed-mode incidence objects required for `FR` and `FM`. The PLY file is downloaded and hash-checked on demand under its stated GNU LGPL license; it is not stored in the repository.

E.3 Urban commuting examples

The urban examples retain the common transport equations but replace labor and income shares with residence and workplace margins. The one-mode oracle first checks the urban welfare derivative in isolation; the multimodal network then adds the shared route and mode machinery.

E.3.1 Three-location regression oracle

The one-mode urban toy follows the residence–workplace structure of Allen and Arkolakis (2022). Its node file records distinct residence and workplace masses, and its edge file records directed commuter traffic. The data are constructed to satisfy the urban flow identity exactly. With spatial externalities and congestion set to zero, the urban derivative collapses to commuter traffic. The test suite also solves the nonlinear exact-hat system around this baseline and compares its central finite difference with the analytical derivative. The example therefore checks the urban Jacobian and welfare projection independently of modal or terminal terms.

E.3.2 Multimodal urban network

The four-location urban example retains the same residence–workplace equilibrium but adds the transport block used by the economic-geography model. Road and transit operate on all ten directed edges, two routes connect the main origin and destination, road traffic creates edge congestion, and transit loads endpoint terminals. Residence shares, workplace shares, OD flows, and edge-mode flows are generated from one contractive route kernel, so no balancing occurs during estimation.

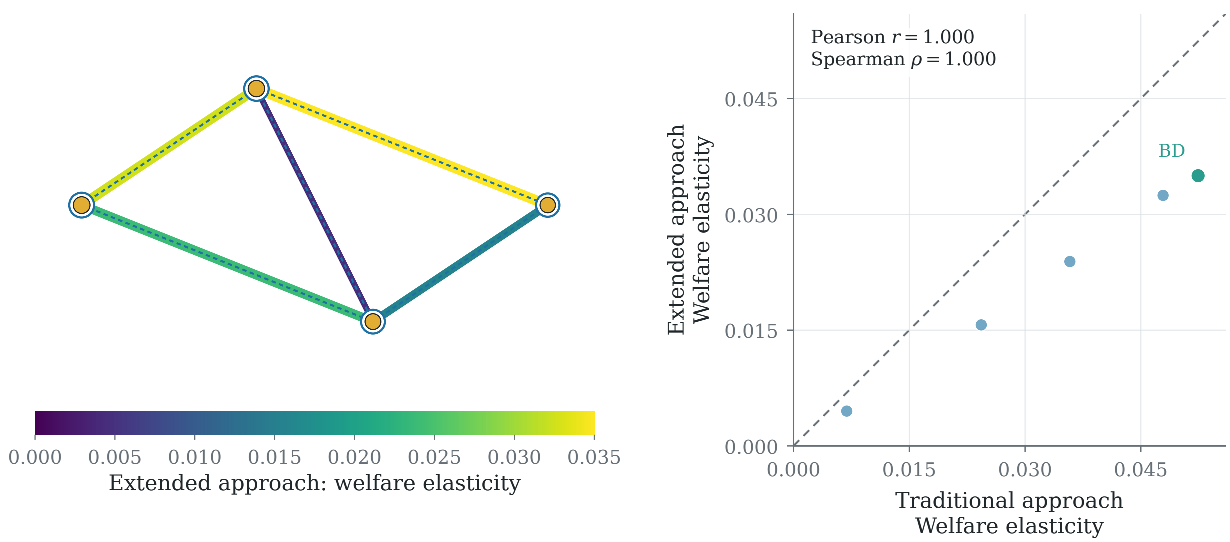

Figure 9: Multimodal urban example. Dashed blue segments identify links with transit service, and blue rings identify terminal locations.

Table 6: Physical links in the multimodal urban example

Link	Traditional	Extended	Traditional rank	Extended rank
BD	0.052417	0.035007	1	1
AB	0.047878	0.032458	2	2
AC	0.035802	0.023888	3	3
CD	0.024323	0.015681	4	4
BC	0.006896	0.004512	5	5

The example changes the spatial residuals and welfare row but reuses the economic-geography route, mode, congestion, and pass-through operators. Its no-congestion derivative equals the urban traffic statistic to within 1×10^{-10}. This is the package's direct numerical check that multimodality belongs to the shared transport block rather than to a separate urban transport model.

E.4 Sioux Falls benchmark adaptation

Sioux Falls is a 24-node test network introduced in early traffic-assignment work and widely used as a small computational benchmark (LeBlanc et al., 1975). The package pins the network, assigned flow, node-coordinate, and trip files in the Transportation Networks for Research repository (Transportation Networks for Research, 2026). Their stated terms limit the files to academic research, so the builder downloads them on demand and verifies the pinned commit and all four hashes.

The benchmark files are not immediately model-ready. OD productions and attractions differ slightly, and the assigned directional flows reproduce those differences. The builder averages production with attraction and each pair of reciprocal assigned flows. The balanced OD margin becomes both labor and income, while balanced link flow is divided by total OD trips. It records the original imbalance and verifies exact conservation after this transformation. The extended configuration evaluates each BPR cost function at the assigned baseline and uses its local log derivative,

$$\lambda_e = \frac{p_e B_e (v_e/c_e)^{p_e}}{1 + B_e (v_e/c_e)^{p_e}},$$

as the link-specific congestion elasticity. This is an adaptation of the assigned baseline to recursive routing, not a new solution of the original Wardrop problem.

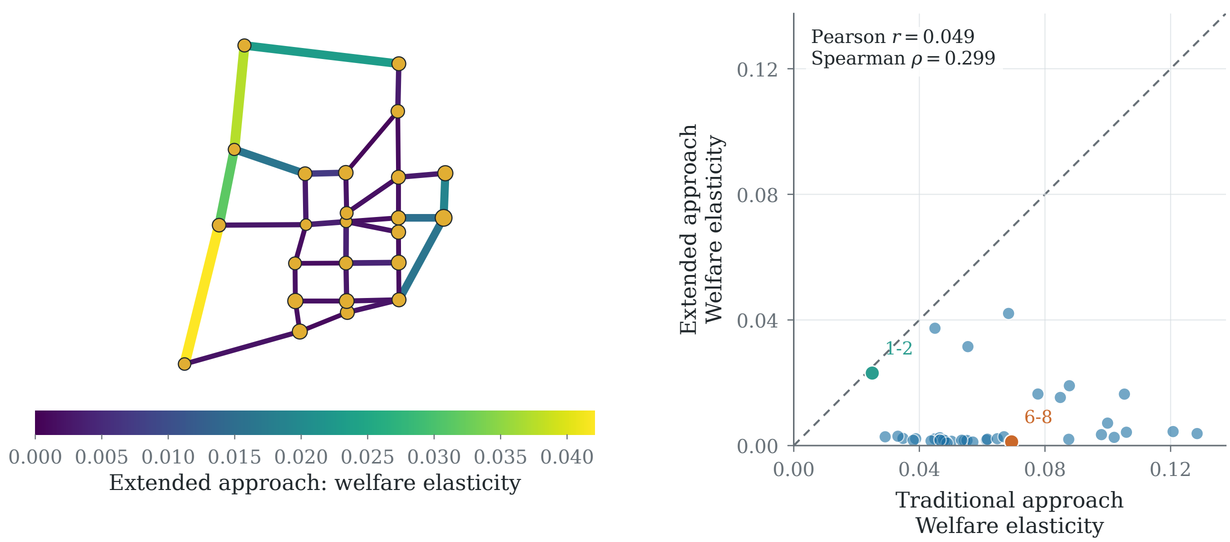

Figure 10: Sioux Falls benchmark adaptation. The calculation uses 24 nodes, 76 directed arcs, and 38 reciprocal physical links.

Table 7: Highest-ranked physical links in the Sioux Falls adaptation

Link	Traditional	Extended	Traditional rank	Extended rank
12-13	0.068403	0.042072	13	1
1-3	0.044963	0.037363	30	2
3-12	0.055452	0.031508	19	3
1-2	0.024996	0.023084	38	4
7-18	0.087767	0.019054	8	5
3-4	0.077751	0.016411	11	6
18-20	0.105295	0.016382	4	7
16-18	0.084891	0.015318	10	8

The Traditional and Extended rankings differ substantially: their Pearson correlation is 0.049, and their rank correlation is 0.299. The mean elasticity falls from 0.064046 to 0.007150. Link-specific BPR derivatives enter the Extended measure but not the Traditional measure, so their variation is one reason the rankings diverge. The calculation is not a validated welfare estimate for Sioux Falls. It shows how the package transforms an assigned-flow traffic benchmark before applying the welfare formula.

E.5 Seattle road diagnostic and multimodal candidate

Seattle provides the larger urban application, but its published road inputs and the candidate multimodal construction have different evidentiary status. The distinction matters before interpreting either set of welfare results.

E.5.1 The published road inputs remain a diagnostic

The first Seattle adapter preserves the 217-location inputs used by Allen and Arkolakis (2022): 2017 LODES residence–workplace flows and 2016 road traffic and travel-time inputs. It verifies the archive hashes, 217 nodes, 1,384 directed road arcs, and 692 reciprocal links. It then fails validation. LODES counts commuters, whereas HPMS annual average daily traffic counts road vehicles from a different population. The two sources imply a maximum normalized recursive-stock disagreement of about 0.108. The package does not silently balance that difference, and the guide therefore reports no welfare map, scatter, or ranking from this rejected baseline.

E.5.2 Construction of the multimodal baseline

The multimodal candidate instead constructs all transport flows from one commuter population. The complete 2017 LODES matrix supplies OD commuter masses and its margins supply residence and workplace activity (U.S. Census Bureau, 2026b). ACS table B08301 supplies each residential origin's public-transit share (U.S. Census Bureau, 2017). A hash-pinned June 2017 King County GTFS feed supplies service, paths, and bus, rail/streetcar, and ferry labels (King County Metro, 2017a). Published Seattle road travel times supply road paths. The builder routes each OD-mode flow over these networks, then checks that the resulting edge-mode traffic reproduces the OD margins.

This construction does not turn schedules into observed ridership. GTFS identifies feasible services and scheduled travel, while ACS identifies mode shares at residential origins. The 2017 Metro System Evaluation provides a separate route-activity comparison and does not determine the OD-to-route allocation (King County Metro, 2017b). The modal elasticity $\eta = 1.099$ is transferred from the economic-geography application and is reported with sensitivity rather than described as a Seattle estimate.

The candidate evaluates 1384 directed road, 627 bus, 12 rail/streetcar, and 1 ferry interventions on the same 217-location equilibrium. Mean Extended-approach gains are 0.000996 percent for bus, 0.002059 percent for rail/streetcar, and 0.001377 percent for ferry. The bus and rail rank correlations with the Traditional approach are 0.954 and 0.972, respectively. The maximum error in the stratified nonlinear finite-difference checks is 9.44×10^{-11}. The accounting identities, deterministic rebuild, and finite-difference check verify the calculation for the stated inputs and calibration. They do not supply

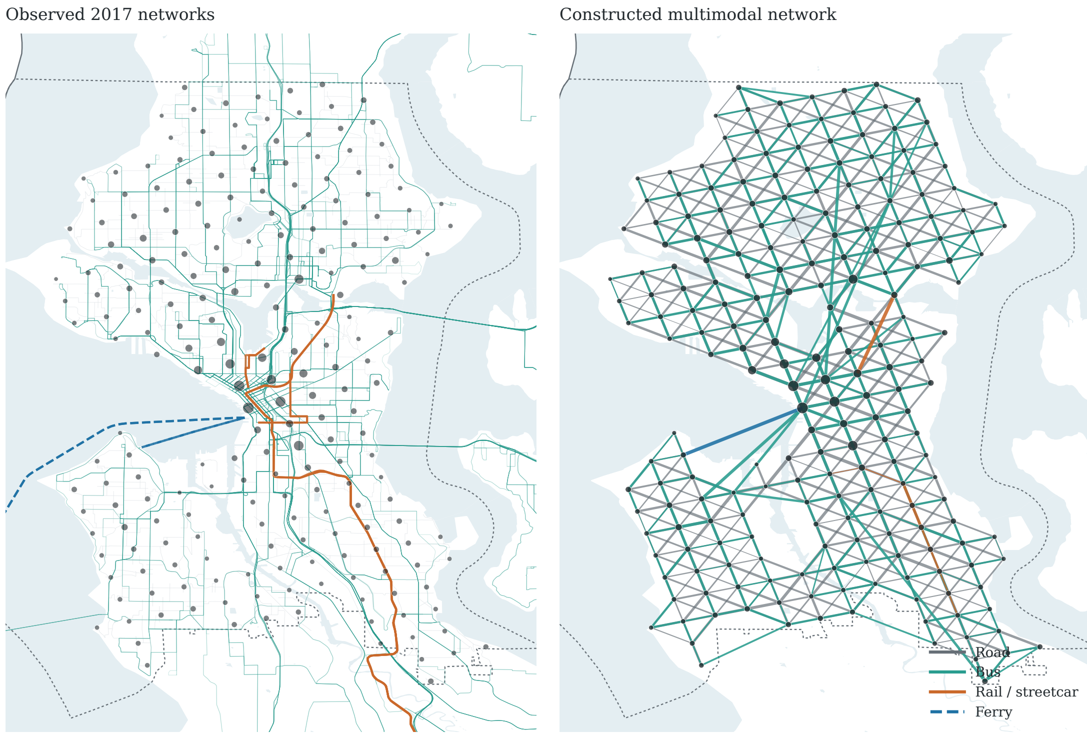

Figure 11: Observed 2017 Seattle networks and the constructed multimodal model network. Road, bus, rail/streetcar, and ferry use distinct colors. Node area records residence plus workplace commuter mass.

Table 8: Selected high-ranked Seattle corridors by mode

Corridor	Mode	Traditional	Extended	Traditional rank	Extended rank
114_121	road	0.139194	0.064713	2	1
100_95	road	0.139007	0.052276	3	2
107_114	road	0.124682	0.050706	4	3
100_95	bus	0.061214	0.024548	1	1
100_106	bus	0.037927	0.021926	2	2
71_95	bus	0.017758	0.020081	10	3
102_123	rail	0.020337	0.012086	1	1
49_67	rail	0.010131	0.007669	2	2
28_49	rail	0.004571	0.002744	3	3
89_95	ferry	0.000888	0.001377	1	1

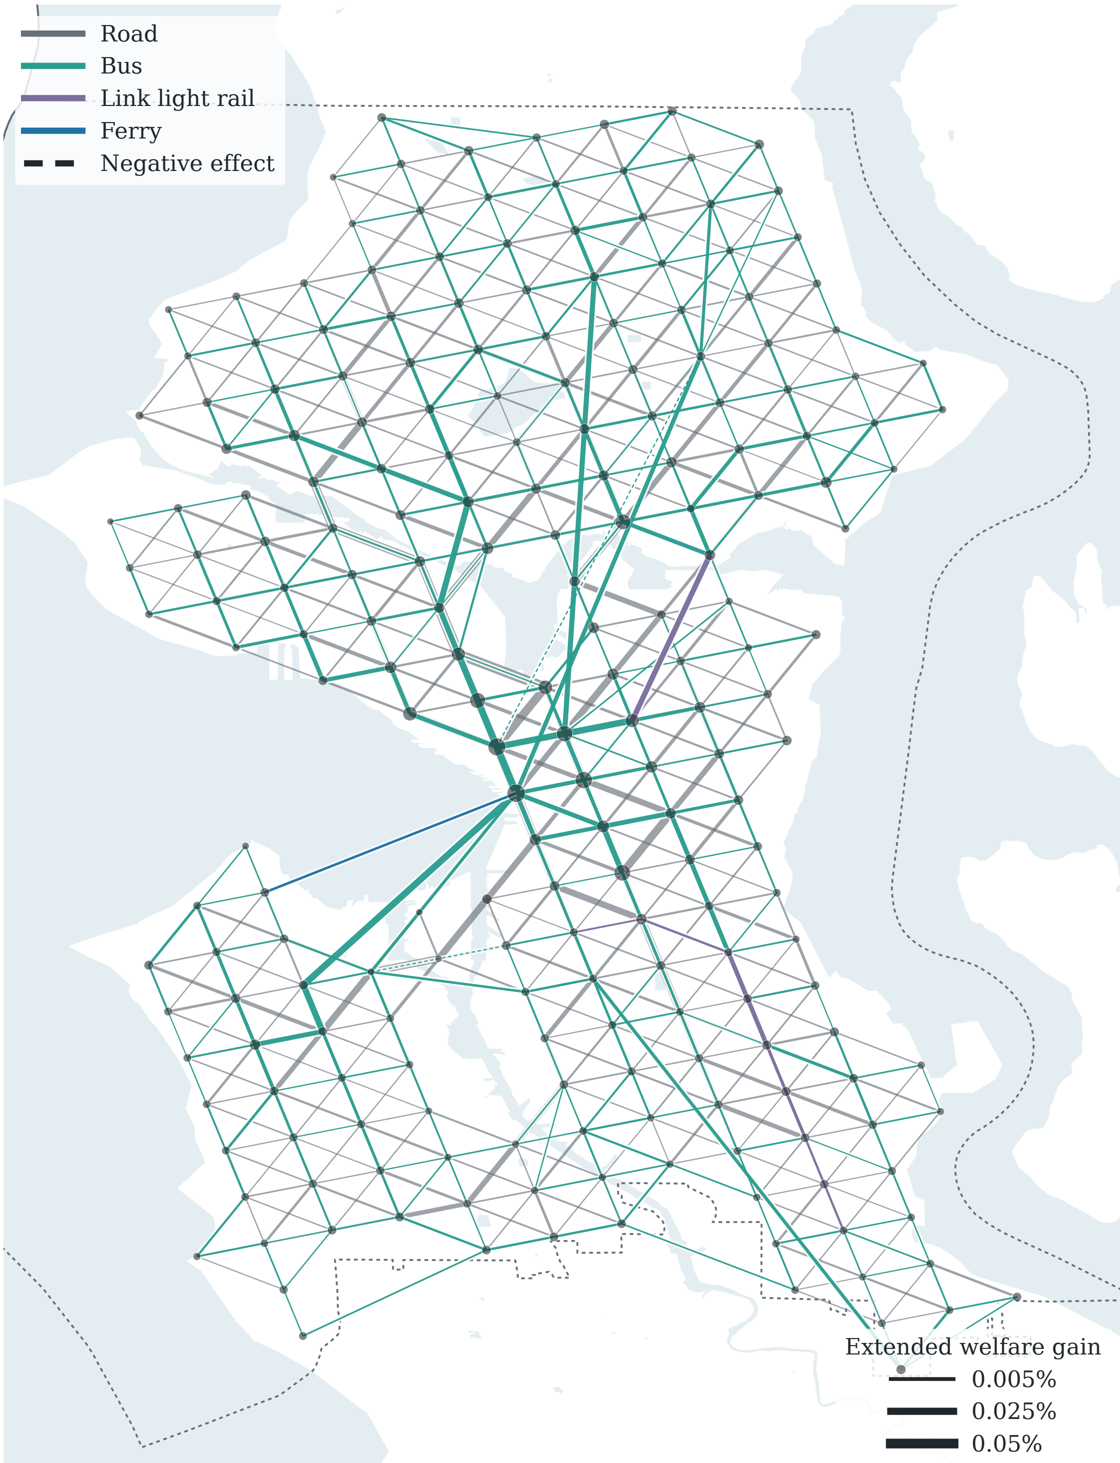

Figure 12: Seattle road and transit policy support. Color identifies the mode, and line width records the Extended-approach welfare gain from a one-percent improvement. Dashed segments have a negative Extended-approach effect.

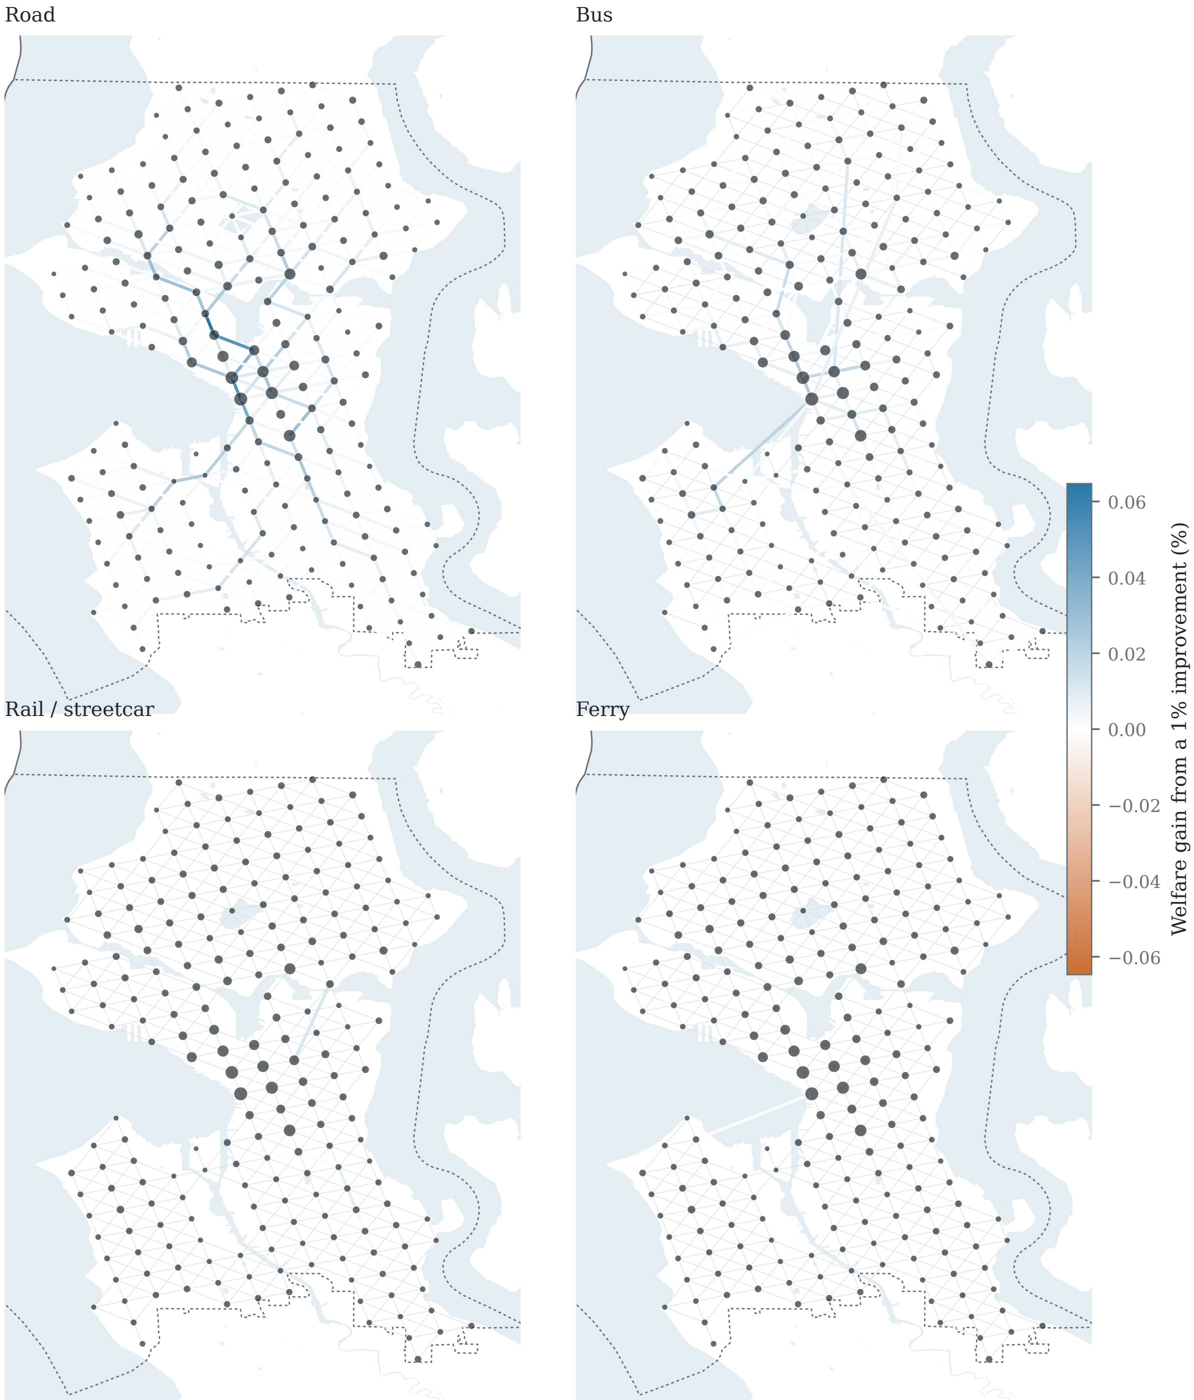

Figure 13: Mode-specific Seattle welfare maps on a common scale. Each panel reports the Extended-approach gain from a one-percent edge-mode improvement, not baseline traffic or ridership.

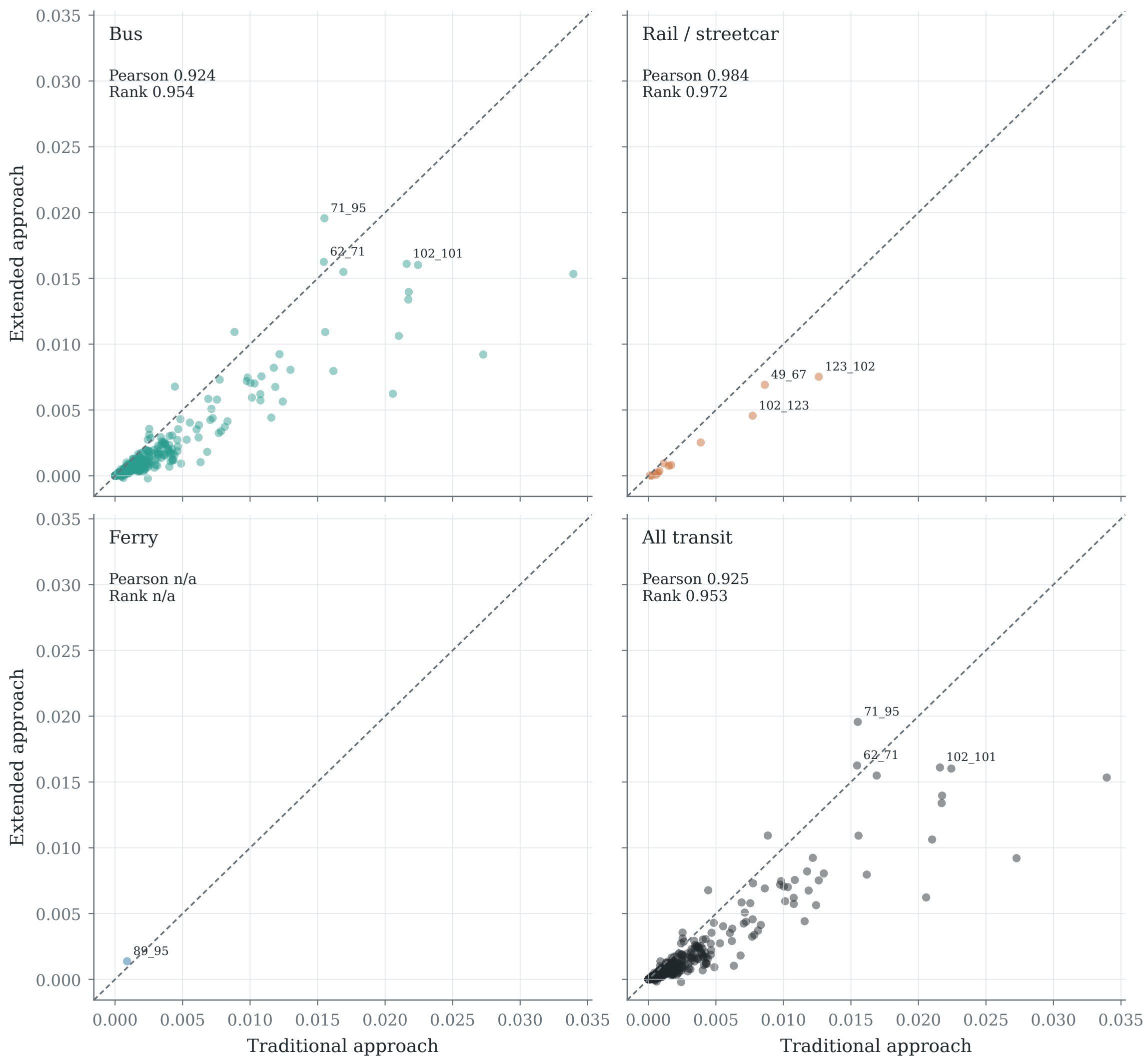

Figure 14: Traditional and Extended approaches for Seattle transit improvements. The four panels share both axes and include the equality line.

route-level ridership, a Seattle estimate of modal substitution, or crowding parameters. The application is therefore an illustration rather than a complete transit appraisal.

E.6 Westeros economic geography

The Westeros example uses the public locations, roads, and continent polygon in CreativeCarto's ArcGIS item (CreativeCarto, 2026). Those layers contain geography but no economic activity or traffic. The current example therefore uses the economic-geography model, not the urban commuting model. The builder hash-checks the source responses, connects road components within a declared 10-kilometer threshold, retains settlements within 50 kilometers of the connected road system, and forms a sparse settlement graph from road distances.

The economic baseline is synthetic and fully disclosed. Location types receive positive common labor and income weights. A symmetric, doubly constrained gravity calculation assigns bilateral trade, and shortest paths route that trade over the reduced road graph. The builder verifies the trade margins and edge-flow accounting before writing the package inputs. The hash-pinned build has 105 locations, 426 directed arcs, and 213 physical links.

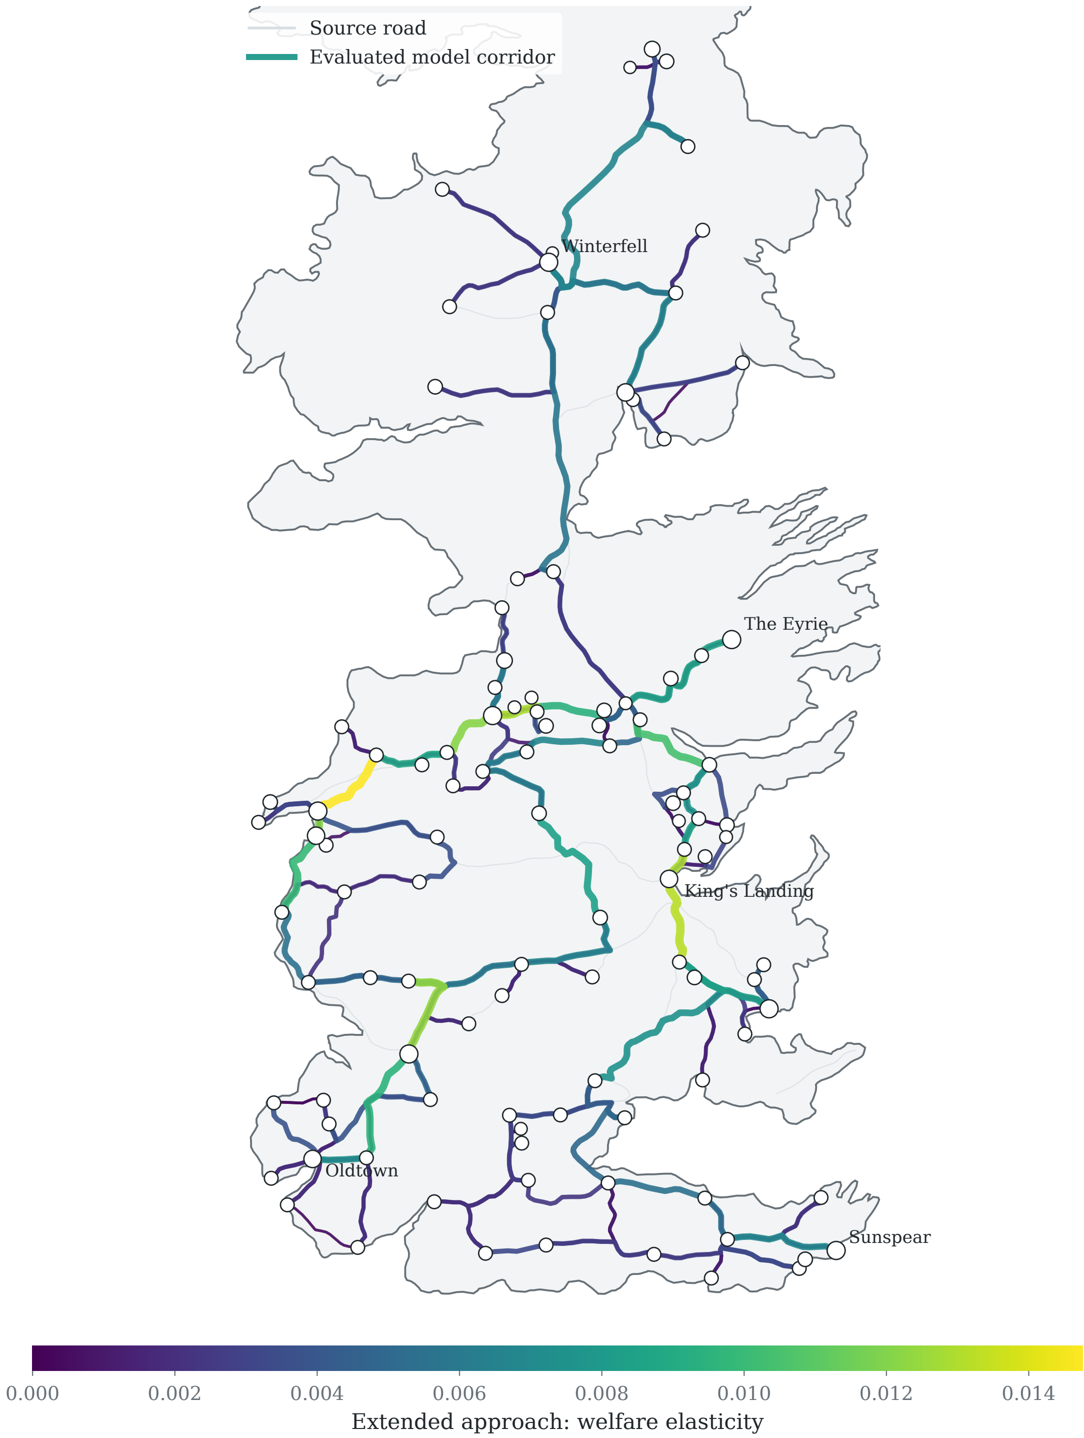

Figure 15: Synthetic Westeros economic geography. The land outline and faint road network come from the CreativeCarto ArcGIS layers (CreativeCarto, 2026). Colored lines trace the evaluated model links along those roads; their color and width record the Extended-approach welfare elasticity. Node area records synthetic economic activity.

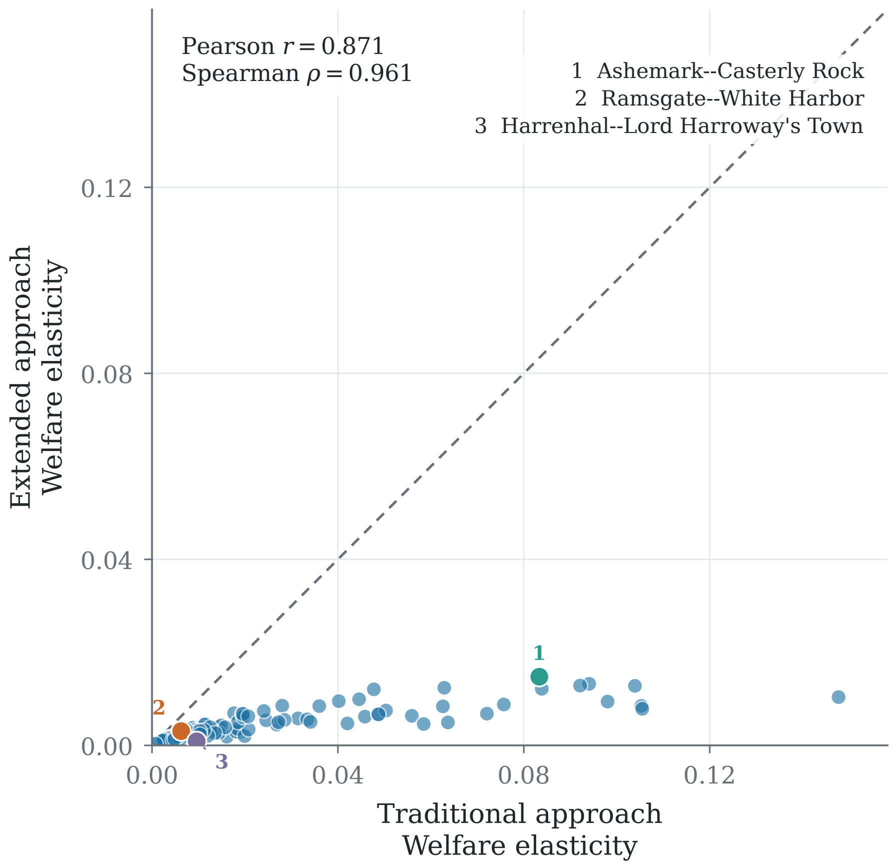

Figure 16: Traditional and Extended welfare elasticities in the synthetic Westeros calculation. Both axes use the same scale, and selected links are identified by their endpoint names.

Table 9: Highest-ranked physical links in the Westeros example

Link	Traditional	Extended	Trad. rank	Ext. rank
Ashemark–Casterly Rock	0.083333	0.014754	9	1
Felwood–King's Landing	0.094063	0.013233	6	2
Riverrun–Stone Hedge	0.092104	0.012854	7	3
Hayford Castle–King's Landing	0.103907	0.012797	4	4
Goldengrove–Highgarden	0.062883	0.012383	13	5
Riverrun–Wayfarer's Rest	0.083861	0.012160	8	6
Casterly Rock–Lannisport	0.047727	0.012054	20	7
Darry–Maidenpool	0.147725	0.010379	1	8

The two welfare measures have a Pearson correlation of 0.871 and a rank correlation of 0.961. Their means are 0.014479 and 0.002747, respectively. These numbers describe how the package behaves on an irregular geography; they do not measure a fictional economy. The ArcGIS item did not state a reuse license when the source manifest was recorded. Raw responses and generated model inputs consequently remain outside Git, and any redistributed display should retain the source attribution.

E.7 Comparison of mechanisms across the self-contained examples

Table 10 reports mean differences across the exact closure cases for the four self-contained examples used in the common asset build. "Congestion" is $\overline{E}^r_F - \overline{E}^r_{NC}$, "Mode flexibility" is $\overline{E}^r_F - \overline{E}^r_{FM}$, and "Route flexibility" is $\overline{E}^r_F - \overline{E}^r_{FR}$. The entries are signed elasticity differences rather than percentage shares, so they may have opposite signs.

Table 10: Mean welfare differences across self-contained examples

Example	Congestion	Mode flexibility	Route flexibility	Top-k overlap
Multimodal grid (EG)	0.001685	-0.000001	-0.000341	5/5
Braess-style routes (EG)	0.012657	0.000000	-0.002924	1/1
Cow network (EG)	0.015092	0.000000	-0.007406	3/5
Multimodal urban (U)	0.002522	-0.000003	-0.000359	1/1

Congestion raises the realized-cost derivative in all four included calibrations, while route flexibility lowers it relative to fixed baseline routes. Modal flexibility is small in these particular synthetic baselines. That pattern is an output, not a maintained restriction. The examples are most useful when read together: the grid combines all channels, the Braess-style topology isolates alternative routes, the cow creates uneven market access on an irregular network, and the urban example changes the spatial equilibrium while retaining the same transport operators.

F Provenance and reproducibility

Reproducing a result requires more than the final CSV. The relevant equations, software, configuration, input vintage, and external sources can have different owners and access conditions. The authority hierarchy below records the governing source for each part of the calculation.

F.1 Authority

The journal manuscript governs paper-facing equations and claims. The package source and tests govern executable behavior. Each run records the package commit, configuration hash, input hashes, model choices, and output hashes. The repository's theory–code and self-containment audits contain the detailed equation mapping, source hierarchy, and kernel lineage; they are not required to build this guide.

The active paper and urban paths use the negative-power `ChoiceLogsum`. `ComponentCES` and positive-power congestion helpers are legacy-only regression paths.

F.2 Synthetic and external example inputs

The six self-contained examples are committed with public synthetic inputs. The common asset target recomputes the worked grid and the three mechanism-focused examples used in the cross-example comparison, then writes a node-and-link visualization contract. The plotting layer turns this contract into welfare maps, comparison scatters, and ranked-link tables; the check target rejects byte-level drift. The smaller API and one-mode urban examples remain verification fixtures and are documented without redundant figures.

Sioux Falls, Seattle, and Westeros require data that are not stored in Git. Their displayed assets were generated from hash-pinned local inputs and are bound to those inputs and outputs in `external-example-assets.json`. A clean clone verifies the frozen asset hashes. A user with the source data can regenerate them with `make practitioner-guide-external-assets`. The Seattle road adapter is not in that asset set because its published commuter and road-traffic inputs fail the package's accounting identity.

F.3 Paper result vintage

Restricted domestic inputs are identified by SHA-256 in the freight replication manifest. The archived result vintage uses the choice logsum, directional Census port flows, edge-local road congestion, and bidirectional physical-link policies. A tracked public-safe aggregate record contains only the statistics reported in this guide and the hashes that bind them to the restricted configuration, inputs, and outputs.

Rebuilding the paper artifacts requires the external files. This PDF includes only aggregate non-sensitive statistics and hashes. It cannot reconstruct the restricted raw data and does not claim to do so.

F.4 Seattle sources

The Seattle candidate depends on a published urban replication archive, LODES, ACS, a pinned historical GTFS feed, and Census geography (Allen and Arkolakis, 2022; U.S. Census Bureau, 2026b, 2017; King County Metro, 2017a). Raw downloads remain outside Git. Their manifests record URLs, versions, retrieval dates, licenses, and hashes. If the exact historical feed is unavailable or a hash changes, the workflow fails rather than substituting current service.

F.5 Guide build

The guide directory contains the PDF source, bibliography, generated TeX, and figures. The PDF build reads only those files and requires no external data or Python process.

The build interface is:

```
make practitioner-guide-assets
make practitioner-guide-example-assets
make practitioner-guide-external-assets
make practitioner-guide
make practitioner-guide-check
```

`practitioner-guide-assets` recomputes all self-contained examples and their figures, tables, and macros. `practitioner-guide-example-assets` redraws the maps, scatters, and ranked-link tables from the already generated visualization contract. `practitioner-guide-external-assets` is optional and requires the hash-pinned Sioux Falls, Seattle, and Westeros roots named in the Makefile. It refreshes only public-safe figures, tables, macros, and their manifest. `practitioner-guide` compiles from the committed assets. `practitioner-guide-check` independently regenerates the source and figure assets for the self-contained examples, verifies the external asset manifest, compiles the PDF, and rejects drift, undefined references, missing figures, private paths, credential-like strings, or a page count outside the stated range.

F.6 Citation and release status

The package will be released under the MIT License following coauthor approval. The preferred paper citation is:

> Allen, Treb, Simon Fuchs, and Woan Foong Wong. "From Social Savings to Spatial Equilibrium: Evaluating Transportation Improvements in Quantitative Spatial Models." 2026.

The software citation should use the tagged release and commit corresponding to the analysis. Feature branches and draft pull requests are not archival releases.

F.7 Reproducibility statement

A fresh user can install the package from a clone, run the committed synthetic examples, reproduce their numerical outputs, and build this guide without private files. Reproducing Sioux Falls, U.S. freight, Seattle, or Westeros results additionally requires the hash-bound external inputs described in their source manifests. The guide is independently buildable; those applications remain versioned external-data calculations.

References

Treb Allen and Costas Arkolakis. Trade and the topography of the spatial economy. *The Quarterly Journal of Economics*, 129(3):1085–1140, 2014. doi: 10.1093/qje/qju016.

Treb Allen and Costas Arkolakis. The welfare effects of transportation infrastructure improvements. Working Paper 25487, National Bureau of Economic Research, 2020. Revised December 2020.

Treb Allen and Costas Arkolakis. The welfare effects of transportation infrastructure improvements. *The Review of Economic Studies*, 89(6):2911–2957, 2022. doi: 10.1093/restud/rdac001.

Treb Allen, Simon Fuchs, and Woan Foong Wong. From social savings to spatial equilibrium: Evaluating transportation improvements in quantitative spatial models. Manuscript, 2026.

David Rezza Baqaee and Emmanuel Farhi. The macroeconomic impact of microeconomic shocks: Beyond hulten's theorem. *Econometrica*, 87(4):1155–1203, 2019. doi: 10.3982/ECTA15202.

Dietrich Braess. Über ein paradoxon aus der verkehrsplanung. *Unternehmensforschung*, 12:258–268, 1968. doi: 10.1007/BF01918335.

John Burkardt. Cow ply dataset. https://people.sc.fsu.edu/~jburkardt/data/ply/cow.ply, 2026. Accessed July 2026.

CreativeCarto. Game of thrones map data layers. https://www.arcgis.com/home/item.html?id=43d03779288048bfb5d3c46e4bc4ccb0, 2026. ArcGIS item; accessed July 2026.

André de Palma, Robin Lindsey, Emile Quinet, and Roger Vickerman, editors. *A Handbook of Transport Economics*. Edward Elgar, Cheltenham, UK, 2011. ISBN 9781847202031. doi: 10.4337/9780857930873.

Dave Donaldson. Transport infrastructure and policy evaluation. In *Handbook of Regional and Urban Economics*, volume 6, pages 287–352. Elsevier, 2025. doi: 10.1016/bs.hesreg.2025.06.004.

Robert William Fogel. A quantitative approach to the study of railroads in american economic growth: A report of some preliminary findings. *The Journal of Economic History*, 22(2):163–197, 1962. doi: 10.1017/S0022050700062719.

Robert William Fogel. *Railroads and American Economic Growth: Essays in Econometric History*. Johns Hopkins Press, Baltimore, 1964.

Simon Fuchs and Woan Foong Wong. Multimodal transport networks. Working Paper 35065, National Bureau of Economic Research, 2026.

Geofabrik. Openstreetmap data extract for greater london. https://download.geofabrik.de/europe/united-kingdom/england/greater-london.html, 2026. OpenStreetMap PBF, Shapefile, and GeoPackage extracts; accessed July 2026.

Charles R. Hulten. Growth accounting with intermediate inputs. *The Review of Economic Studies*, 45(3):511–518, 1978. doi: 10.2307/2297252.

King County Metro. General transit feed specification, june 2017 service. https://www.transit.land/feeds/f-c23-metrokingcounty/versions/ad172e653aa881557a5f3cb84f2ace6819308600, 2017a. Feed-version SHA-1 ad172e653aa881557a5f3cb84f2ace6819308600.

King County Metro. 2017 system evaluation. Technical report, King County Department of Transportation, 2017b. URL https://cdn.kingcounty.gov/-/media/king-county/depts/metro/documents/about/data-and-reports/2017/2017-system-evaluation.pdf.

Larry J. LeBlanc, Edward K. Morlok, and William P. Pierskalla. An efficient approach to solving the road network equilibrium traffic assignment problem. *Transportation Research*, 9(5):309–318, 1975. doi: 10.1016/0041-1647(75)90030-1.

Office for National Statistics. Middle layer super output areas, december 2021, boundaries for england and wales. https://geoportal.statistics.gov.uk/, 2021. Open Geography Portal; accessed July 2026.

Office for National Statistics. Travel to work quality information for census 2021. https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/methodologies/traveltoworkqualityinformationforcensus2021, 2022. Accessed July 2026.

Office for National Statistics. Origin-destination data, england and wales: Census 2021. https://www.nomisweb.co.uk/sources/census_2021_od, 2023. Dataset ODWP01EW and accompanying metadata; accessed July 2026.

Office of Rail and Road. Estimates of station usage, april 2023 to march 2024. https://dataportal.orr.gov.uk/station-usage, 2024. Official statistics and origin–destination data; accessed July 2026.

Stephen J. Redding and Matthew A. Turner. Transportation costs and the spatial organization of economic activity. In *Handbook of Regional and Urban Economics*, volume 5, pages 1339–1398. Elsevier, 2015. doi: 10.1016/B978-0-444-59531-7.00020-X.

Kenneth A. Small, Erik T. Verhoef, and Robin Lindsey. *The Economics of Urban Transportation*. Routledge, London, 3 edition, 2024. ISBN 9781138069053. doi: 10.4324/9781315157375.

Transport for London. Network demand data. https://tfl.gov.uk/corporate/publications-and-reports/network-demand-data, 2026a. Ticketing-based journey and station-footfall data; accessed July 2026.

Transport for London. Our open data. https://tfl.gov.uk/info-for/open-data-users/our-open-data, 2026b. Journey-planning, timetable, and station-location feeds; accessed July 2026.

Transportation Networks for Research. Sioux falls transportation network. https://github.com/bstabler/TransportationNetworks/tree/977ee75c6906337c0c7d229a1336107c7cdb533e/SiouxFalls, 2026. GitHub repository pinned at commit 977ee75c6906337c0c7d229a1336107c7cdb533e; accessed July 2026.

U.S. Census Bureau. 2017 american community survey five-year estimates, table b08301: Means of transportation to work. https://api.census.gov/data/2017/acs/acs5/groups/B08301.html, 2017.

U.S. Census Bureau. International trade time series: Port-level exports and imports. https://www.census.gov/data/developers/data-sets/international-trade.html, 2026a. Port-level API data for 2017; accessed July 2026.

U.S. Census Bureau. Longitudinal employer–household dynamics: Origin–destination employment statistics data. https://lehd.ces.census.gov/data/, 2026b. Accessed July 2026.