

Quantitative Spatial Economics

The Allen–Arkolakis Model Calibrated to Italy

Technical Documentation and Results

EUI Economics Course

Abstract

This document describes the MATLAB implementation of the [Allen and Arkolakis \(2014\)](#) quantitative spatial model, calibrated to the 20 administrative regions of Italy. The code solves for the spatial equilibrium, performs a double model inversion to recover both regional productivities and amenities, and runs a series of counterfactual policy experiments that illustrate the core mechanisms of the model: agglomeration, congestion, trade costs, and labor mobility. Seven figures are produced from the main model, covering the data, the baseline fundamentals, bilateral trade, agglomeration–congestion comparative statics, a Mezzogiorno productivity subsidy, high-speed rail investment, and a comparison of hub-and-spoke versus ring infrastructure. A companion hat algebra implementation validates the exact and linear counterfactual methods from Section 6 of [Allen and Arkolakis \(2025\)](#), producing three additional figures. A second companion script implements the estimation and model-testing procedures from Section 7 of [Allen and Arkolakis \(2025\)](#): gravity estimation of trade costs via PPML, network-based trade frictions, estimation of model elasticities via both reduced-form regressions and a minimum distance estimator that embeds the exact hat algebra inside the optimization loop, the market access approach to counterfactuals, and ex-post model testing, producing four additional figures.

Contents

1	Introduction	3
2	Data	3
2.1	Regional Data	3
2.2	Region Boundaries	4
2.3	Data Limitations	4
3	The Model	4
3.1	Preferences and Trade	5
3.2	Production and Agglomeration	5
3.3	Amenities and Congestion	5
3.4	Spatial Equilibrium	5
4	Calibration Strategy	6
4.1	Parameters Set Directly	7
4.2	Trade Costs	7
4.3	Innate Productivity	7
5	Solution Algorithm	8
6	Double Model Inversion	8

7	Results and Experiments	9
7.1	Results	10
7.1.1	Baseline Equilibrium (Figure 2)	10
7.1.2	Bilateral Trade (Figure 3)	11
7.1.3	Agglomeration and Congestion (Figure 4)	11
7.2	Policy Experiments	12
7.2.1	Mezzogiorno Productivity Subsidy (Figure 5)	12
7.2.2	High-Speed Rail Corridor (Figure 6)	13
7.2.3	Hub-and-Spoke vs. Southern Ring (Figure 7)	14
8	Hat Algebra: Counterfactual Analysis Without Fundamentals	15
8.1	Notation and Setup	16
8.2	Elasticities and Uniqueness	16
8.3	Exact Hat Algebra (Nonlinear Iteration)	16
8.4	Linear Comparative Statics	17
8.5	Experiments and Validation	17
8.5.1	Experiment 1: HSR Corridor (30% Reduction)	17
8.5.2	Experiment 2: Mezzogiorno Subsidy (20%)	18
8.5.3	Accuracy vs. Shock Size	19
9	Bringing the Model to the Data	20
9.1	Gravity Estimation of Trade Costs	20
9.2	Network-Based Trade Costs	22
9.3	Estimation of Model Elasticities	23
9.4	The Market Access Approach	25
9.5	Testing the Model	26
10	Interpreting the Results	28
11	Technical Notes	28
11.1	Software Requirements	28
11.2	Computational Performance	28
11.3	Output Files	28
11.4	Extending the Code	29

1 Introduction

The Allen–Arkolakis framework (Allen and Arkolakis, 2014) is an Armington trade model with free labor mobility, productivity agglomeration, and amenity congestion. It nests several classic spatial models—including the Rosen–Roback model and elements of Krugman (1991)—and admits a clean characterization of equilibrium. When agglomeration and congestion elasticities are zero, the equilibrium can be solved in closed form via the Perron–Frobenius theorem. For the general case (nonzero elasticities), the code uses an iterative fixed-point algorithm.

The Italian calibration exploits the sharp North–South productivity gradient and the country’s distinctive geography (including two major islands) to produce experiments that are economically meaningful and pedagogically rich. The model is calibrated to approximate 2019 data from ISTAT and Eurostat, with 20 regions corresponding to Italy’s NUTS-2 administrative divisions.

The main code (`solve_aa_italy.m`) is self-contained: it requires no MATLAB toolboxes and produces seven PNG figures at 200 dpi along with a summary table printed to the command window. A companion script (`solve_aa_italy_hatalgebra.m`) implements the hat algebra from Section 6 of Allen and Arkolakis (2025), comparing exact nonlinear, linear, and full re-solve methods across the same Italian geography.

2 Data

2.1 Regional Data

All data are approximate 2019 values. Table 1 summarizes the inputs for each region.

Table 1: Italian Regional Data (approx. 2019, ISTAT / Eurostat / INPS)

#	Region	Code	Macro	Pop. (M)	GDP/cap (k€)	Wage (€)	Lat	Lon
1	Piemonte	PIE	North	4.36	31.0	31 448	45.07	7.69
2	Valle d’Aosta	VDA	North	0.13	38.0	31 128	45.74	7.32
3	Lombardia	LOM	North	10.06	38.0	33 452	45.47	9.19
4	Trentino-A.A.	TAA	North	1.07	41.0	31 706	46.50	11.35
5	Veneto	VEN	North	4.91	32.0	30 848	45.44	11.99
6	Friuli V.G.	FVG	North	1.22	30.0	30 872	46.07	13.23
7	Liguria	LIG	North	1.55	30.0	32 156	44.41	8.95
8	Emilia-Romagna	EMR	North	4.46	35.0	31 441	44.49	11.34
9	Toscana	TOS	Center	3.73	30.0	29 884	43.35	11.17
10	Umbria	UMB	Center	0.88	24.0	28 530	42.96	12.39
11	Marche	MAR	Center	1.53	26.0	28 852	43.37	13.18
12	Lazio	LAZ	Center	5.88	33.0	32 360	41.90	12.49
13	Abruzzo	ABR	South	1.31	24.0	28 641	42.35	13.39
14	Molise	MOL	South	0.31	20.0	27 263	41.56	14.66
15	Campania	CAM	South	5.80	18.0	27 606	40.83	14.25
16	Puglia	PUG	South	4.03	18.0	27 261	41.13	16.87
17	Basilicata	BAS	South	0.56	20.0	26 055	40.64	15.80
18	Calabria	CAL	South	1.95	16.5	26 631	38.91	16.59
19	Sicilia	SIC	South	5.00	17.0	27 289	37.60	14.02
20	Sardegna	SAR	South	1.64	21.0	27 294	39.23	9.12

Population is regional resident population in millions (total ≈ 60.4 M), normalized to shares summing to 1 in the model. **GDP per capita** in thousands of euros serves as the initial proxy for regional productivity in the double inversion. **Wage** is the average annual gross salary from INPS administrative records; normalized to mean 1, this is the wage target that the model

matches exactly via the double inversion. The wage gradient is flatter than the GDP/cap gradient: Northern wages range 30.8–33.5 k€ while Southern wages range 26.1–28.6 k€ (a ratio of about 1.2:1, compared to nearly 2:1 for GDP/cap). This difference reflects the fact that GDP per capita conflates wages with capital income, transfers, and employment rate variation. **Coordinates** are approximate geographic centroids used for computing bilateral distances.

Figure 1 displays the two key observables that the model will target: equilibrium wages and population shares across the 20 regions.

Figure 1: Data -- Wages and Population Shares

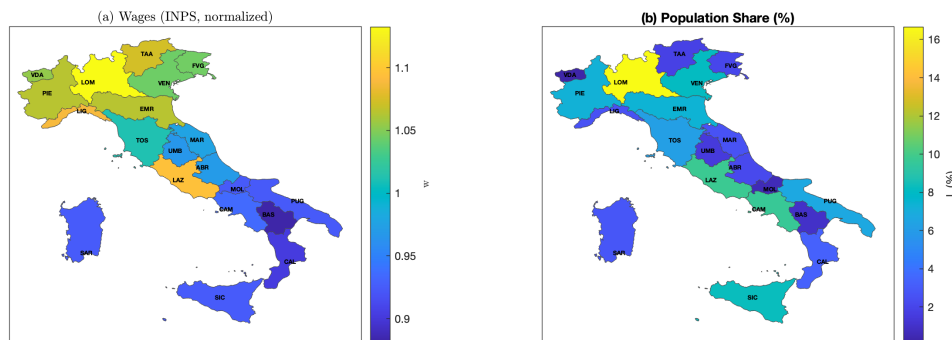


Figure 1: Data. (a) Average annual wages (INPS, 2019) by region; (b) population shares by region.

2.2 Region Boundaries

Choropleth maps use an ESRI shapefile of Italian regional boundaries (source: SimpleMaps), pre-converted to a MATLAB .mat file (`italy_regions.mat`) to avoid requiring the Mapping Toolbox. Each region is identified by its ISTAT code and mapped to the internal ordering. Multipart polygons (e.g., Sicilia’s smaller islands) are handled with NaN separators.

2.3 Data Limitations

Several simplifications should be noted:

- GDP per capita is used only as the *initial guess* for innate productivity; the final \bar{A}_j are recovered by the double inversion. GDP per capita conflates productivity with capital income, transfers, and employment rates, but the inversion corrects for this.
- Average annual wages (INPS) are the target the model matches. These are gross salaries from administrative records and are a cleaner measure of labor compensation than GDP/cap, though they still reflect composition effects (sectoral mix, hours, skill distribution).
- Centroids are geographic rather than population-weighted, which matters for large regions where population is concentrated in one area (e.g., eastern Sicilia).
- The data are rounded to the nearest 0.5 k€ for GDP per capita and 0.01 M for population.

These are standard approximations in teaching implementations and do not affect the qualitative conclusions.

3 The Model

The model follows Section 4 of Allen and Arkolakis (2014). There are $N = 20$ locations (regions) and a total labor supply $\bar{L} = 1$ (normalized). Workers are freely mobile across regions.

3.1 Preferences and Trade

Each region j produces a single differentiated variety. Consumers in region j have CES preferences over varieties from all regions with elasticity of substitution $\sigma > 1$. The price of region i 's good in region j is

$$p_{ij} = \frac{\tau_{ij} w_i}{A_i},$$

where $\tau_{ij} \geq 1$ is the iceberg trade cost, w_i is the wage, and A_i is effective (total) productivity. The CES price index in region j is

$$P_j = \left[\sum_{i=1}^N \left(\frac{\tau_{ij} w_i}{A_i} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}. \quad (1)$$

Bilateral trade shares follow the standard gravity form:

$$\pi_{ij} = \frac{(\tau_{ij} w_i / A_i)^{1-\sigma}}{P_j^{1-\sigma}}. \quad (2)$$

3.2 Production and Agglomeration

Each region has innate productivity \bar{A}_j (a fundamental). Effective productivity depends on local employment through an agglomeration externality:

$$A_j = \bar{A}_j \cdot L_j^\alpha, \quad (3)$$

with $\alpha \geq 0$. When $\alpha > 0$, larger labor forces generate productivity spillovers (knowledge sharing, thicker intermediate input markets, better labor matching). In the calibration, $\alpha = 0.05$, representing mild agglomeration.

3.3 Amenities and Congestion

Each region has innate amenity \bar{u}_j (also a fundamental). Effective amenity depends on local population through a congestion externality:

$$u_j = \bar{u}_j \cdot L_j^\beta, \quad (4)$$

with $\beta \leq 0$. When $\beta < 0$, larger populations reduce quality of life (housing costs, commuting, pollution, crowding). In the calibration, $\beta = -0.20$, representing moderate congestion.

3.4 Spatial Equilibrium

Given a geography $\{\bar{A}_j, \bar{u}_j, \tau_{ij}\}_{i,j=1}^N$, a *spatial equilibrium* is a vector of wages $\{w_j\}$, population shares $\{L_j\}$, and an aggregate welfare level W such that:

(i) *Price indices*. The CES price index in each region is consistent with wages and productivities:

$$P_j = \left[\sum_{i=1}^N \left(\frac{\tau_{ij} w_i}{\bar{A}_i L_i^\alpha} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}. \quad (5)$$

(ii) *Goods market clearing*. Each region's total revenue equals its wage bill:

$$w_j L_j = \sum_{i=1}^N \pi_{ji} w_i L_i, \quad (6)$$

where the trade shares are $\pi_{ji} = (\tau_{ji} w_j / A_j)^{1-\sigma} / P_i^{1-\sigma}$.

(iii) *Free mobility*. Workers equalize indirect utility across all populated regions:

$$V_j = \frac{w_j \cdot \bar{u}_j \cdot L_j^\beta}{P_j} = W \quad \text{for all } j \text{ with } L_j > 0. \quad (7)$$

(iv) *Labor market clearing*. Total population is fixed:

$$\sum_{j=1}^N L_j = \bar{L}. \quad (8)$$

Conditions (6) and (7) form a system of $2N$ equations in $2N$ unknowns (wages and populations), with W determined residually.

Combined equilibrium conditions. Substituting the CES price index, trade shares, agglomeration, and congestion into (6) and (7) yields the two key equations of the model (equations 31 and 32 of Allen and Arkolakis 2014):

$$W^{\sigma-1} w_i^\sigma L_i^{1-\alpha(\sigma-1)} = \sum_{j=1}^N \left(\frac{\bar{A}_i \bar{u}_j}{\tau_{ij}} \right)^{\sigma-1} w_j^\sigma L_j^{1+\beta(\sigma-1)}, \quad (31)$$

$$W^{\sigma-1} w_i^{1-\sigma} L_i^{\beta(1-\sigma)} = \sum_{j=1}^N \left(\frac{\bar{u}_i \bar{A}_j}{\tau_{ji}} \right)^{\sigma-1} w_j^{1-\sigma} L_j^{\alpha(\sigma-1)}. \quad (32)$$

Equation (31) captures total revenue: the left-hand side is region i 's wage bill (scaled by welfare and agglomeration), and the right-hand side sums the revenue i earns from selling to every destination j . Equation (32) captures total expenditure: the left-hand side is region i 's spending power (scaled by welfare and congestion), and the right-hand side sums the expenditure i faces importing from every origin j . Together with the labor-clearing constraint $\sum_i L_i = \bar{L}$, these form a system of $2N + 1$ equations in $2N + 1$ unknowns ($\{w_i, L_i\}_{i=1}^N$ and W).

For the special case $\alpha = \beta = 0$, the equilibrium can be solved in closed form via the Perron–Frobenius theorem (Allen and Arkolakis, 2014). For general elasticities, the equilibrium is computed numerically (Section 5).

4 Calibration Strategy

The calibration follows the standard model-inversion approach in quantitative spatial economics (Redding and Rossi-Hansberg, 2017).

4.1 Parameters Set Directly

Table 2: Model Parameters

Parameter	Value	Interpretation
σ	5	Elasticity of substitution. Standard in the trade literature.
α	0.05	Agglomeration: $A_j = \bar{A}_j L_j^\alpha$. Mild positive spillovers from density.
β	-0.20	Congestion: $u_j = \bar{u}_j L_j^\beta$. Crowding, pollution, and housing costs.
ρ	1.0	Distance elasticity (per 1000 km). Trade cost: $\tau_{ij} = e^{\rho d_{ij}/1000}$.
Sardegna	+20%	Multiplicative surcharge on all trade costs (island isolation).
Sicilia	+8%	Smaller surcharge (Strait of Messina).

The value $\sigma = 5$ is standard in the international/regional trade literature. The agglomeration and congestion elasticities are within the range of empirical estimates surveyed in [Combes and Gobillon \(2015\)](#). The distance elasticity $\rho = 1.0$ is chosen so that the model generates plausible internal trade shares.

4.2 Trade Costs

Bilateral trade costs are parameterized as:

$$\tau_{ij} = \exp\left(\rho \cdot \frac{d_{ij}}{1000}\right), \quad (9)$$

where d_{ij} is the geographic distance in kilometers between region centroids and $\rho = 1.0$ is the distance elasticity per 1000 km. Distances are computed using a planar approximation (1° latitude ≈ 111 km, 1° longitude ≈ 85 km at 42° N). Internal trade costs are $\tau_{ii} = 1$.

The effective trade elasticity is $(\sigma - 1) \times \rho = 4.0$ per 1000 km. For two regions 400 km apart (e.g., Lombardia and Toscana), $\tau^{\sigma-1} = e^{0.16 \times 4} \approx 1.90$, producing a moderate penalty relative to the home link. This generates home trade shares in the range of 20–40%, with substantial bilateral trade flows.

Two adjustments capture island isolation:

- **Sardegna:** all bilateral trade costs multiplied by 1.20 (a 20% sea-crossing penalty reflecting the lack of a fixed link and longer shipping routes).
- **Sicilia:** a smaller 8% penalty (the Strait of Messina is narrow, and ferry connections are frequent).

4.3 Innate Productivity

The initial guess for innate productivity \bar{A}_j is set proportional to GDP per capita, normalized to have a mean of 1 across regions:

$$\bar{A}_j^{(0)} = \frac{\text{GDP}/\text{cap}_j}{\text{GDP}/\text{cap}}$$

This is only a starting point: GDP per capita reflects equilibrium wages (which depend on \bar{A} , L , and trade), not just productivity fundamentals. The final values of \bar{A}_j are recovered jointly with amenities via the double inversion (Section 6).

5 Solution Algorithm

Given a set of fundamentals $(\bar{A}_j, \bar{u}_j, \tau_{ij})$, the equilibrium defined in Section 3.4 is solved iteratively using the combined equilibrium conditions (31)–(32).

The ratio trick. Dividing (31) by (32) eliminates the welfare level W . Collecting all w_i terms on the left-hand side gives $w_i^{2\sigma-1}$, and the population terms contribute $L_i^{1+(\beta-\alpha)(\sigma-1)}$. Solving for w_i :

$$w_i^{\text{new}} = \left(\frac{\text{RHS of (31)}}{\text{RHS of (32)}} \cdot L_i^{-\gamma_L} \right)^{1/\gamma_w},$$

where $\gamma_L = 1 + (\beta - \alpha)(\sigma - 1)$ and $\gamma_w = 2\sigma - 1$. Normalize to mean 1.

Population update. Given updated wages, use (31) alone to recover the implied population:

$$Z_i = \left(\frac{\text{RHS of (31)}}{(w_i^{\text{new}})^\sigma} \right)^{1/(1-\alpha(\sigma-1))},$$

then set $L_i^{\text{new}} = Z_i / \sum_k Z_k \cdot \bar{L}$.

Damping. To ensure stability, wages and populations are updated via log-linear blending with factor $d = 0.1$:

$$\log L_i^{t+1} = (1 - d) \log L_i^t + d \log L_i^{\text{new}},$$

and similarly for wages.

Convergence is declared when $\max_i (|\Delta \log L_i|, |\Delta \log w_i|) < 10^{-10}$. The algorithm typically converges in a few hundred iterations for the baseline, and in up to 15 000 iterations for more extreme parameter configurations. The damping factor of 0.1 is conservative but ensures stability.

6 Double Model Inversion

Neither innate productivity \bar{A}_j nor innate amenities \bar{u}_j can be directly observed. The goal of the inversion is to recover both sets of fundamentals so that the model's equilibrium exactly reproduces the observed wages w_j^{data} and population shares L_j^{data} . The standard approach follows Redding and Rossi-Hansberg (2017): with N regions, the model has $2N$ unobserved fundamentals (\bar{A}_j and \bar{u}_j) and the data provide $2N$ targets (N population shares and N wages from INPS administrative records, normalized to mean 1). The system is exactly identified.

Initial guesses. The inversion requires starting values:

1. **Productivity.** Set $\bar{A}_j^{(0)}$ proportional to GDP per capita (normalized to mean 1).
2. **Amenities.** Given $\bar{A}^{(0)}$, w^{data} , L^{data} , and τ , compute price indices P_j from (5), then invert the free-mobility condition (7):

$$\bar{u}_j^{(0)} = \frac{W \cdot P_j}{w_j^{\text{data}} \cdot (L_j^{\text{data}})^{\beta'}}$$

with $W = 1$ as normalization, rescaled to mean 1.

These one-shot approximations use data-implied wages and price indices rather than self-consistent equilibrium objects, so the model solved with $(\bar{A}^{(0)}, \bar{u}^{(0)})$ will not exactly reproduce the data.

Iterative double inversion. We embed the solution algorithm of Section 5 in an outer loop that jointly updates \bar{A} and \bar{u} until both $L^{\text{model}} = L^{\text{data}}$ and $w^{\text{model}} = w^{\text{data}}$:

1. Start from $(\bar{A}^{(0)}, \bar{u}^{(0)})$.
2. At each outer iteration k :
 - (a) Solve the full GE model with current $(\bar{A}^{(k)}, \bar{u}^{(k)})$ using the iterative solver, obtaining $(L^{(k)}, w^{(k)})$.
 - (b) If $\max_j (|L_j^{(k)} - L_j^{\text{data}}|, |w_j^{(k)} - w_j^{\text{data}}|) < \text{tol}$, stop.
 - (c) Otherwise, apply log-linear corrections:

$$\log \bar{u}_j^{(k+1)} = \log \bar{u}_j^{(k)} + \gamma_u \cdot \log \left(\frac{L_j^{\text{data}}}{L_j^{(k)}} \right),$$

$$\log \bar{A}_j^{(k+1)} = \log \bar{A}_j^{(k)} + \gamma_A \cdot \log \left(\frac{w_j^{\text{data}}}{w_j^{(k)}} \right).$$

Re-normalize both to mean 1.

The intuition is direct: if a region has too few people, raise its amenity; if its wages are too low, raise its productivity. The damping factors $\gamma_u = 0.3$ and $\gamma_A = 0.2$ (productivity updates are more cautious because of stronger cross-effects through trade) ensure stable convergence.

Why double inversion matters. If only amenities are inverted—holding \bar{A}_j fixed at the GDP/cap proxy—the model exactly matches populations but *not* wages. The resulting baseline does not faithfully represent the current economy, so counterfactual predictions depart from the wrong starting point. The double inversion ensures that the baseline is an exact representation of the data, which is the methodologically correct foundation for credible policy analysis.

7 Results and Experiments

The code produces seven figures, each saved as a PNG file at 200 dpi. Table 3 summarizes the outputs.

Table 3: Summary of Figures

Fig.	Content	Description
1	Data maps	Choropleth maps of wages and population shares (observed data).
2	Baseline equilibrium	Model fit (population and wage scatters) and backed-out fundamentals \bar{A}_j, \bar{u}_j .
3	Bilateral trade	Trade share matrix π_{ij} (non-targeted prediction).
4	Agglom./congestion	Comparative statics: L_j for 4 representative regions and Gini vs. α and β .
5	Mezzogiorno subsidy	Place-based policy: raise \bar{A} in 8 Southern regions by 0–20%. Maps at 5%, welfare path.
6	HSR corridor	Reduce τ_{ij} along Milano–Roma–Napoli by 0–60%. Choropleth with HSR overlay, welfare path.
7	Hub vs. Ring	Equal-budget: reduce all Roma routes by 25% vs. connect 8 Southern regions in a ring.

7.1 Results

7.1.1 Baseline Equilibrium (Figure 2)

Figure 2 summarizes the baseline equilibrium. Panels (a) and (b) confirm that the double inversion achieves a perfect fit: model population shares and wages lie exactly on the 45-degree line, with correlations of 1.0000. This is by construction—the inversion recovers both \bar{A}_j and \bar{u}_j to match the data—and ensures that all subsequent counterfactuals depart from the correct starting point.

Panels (c) and (d) display the backed-out fundamentals. Innate productivity \bar{A}_j broadly follows the North–South GDP gradient, but the inverted values differ from the raw GDP/cap proxy because the inversion strips out the contribution of agglomeration spillovers and trade access. The backed-out amenities \bar{u}_j exhibit the *opposite* pattern: Southern regions have high amenities, Northern regions have low amenities. This is the model’s rationalization of why the South is populated despite low wages—given their low productivity and high prices, there must be compensating amenities (climate, cultural heritage, family networks, lower cost of non-traded goods) that keep people from moving North.

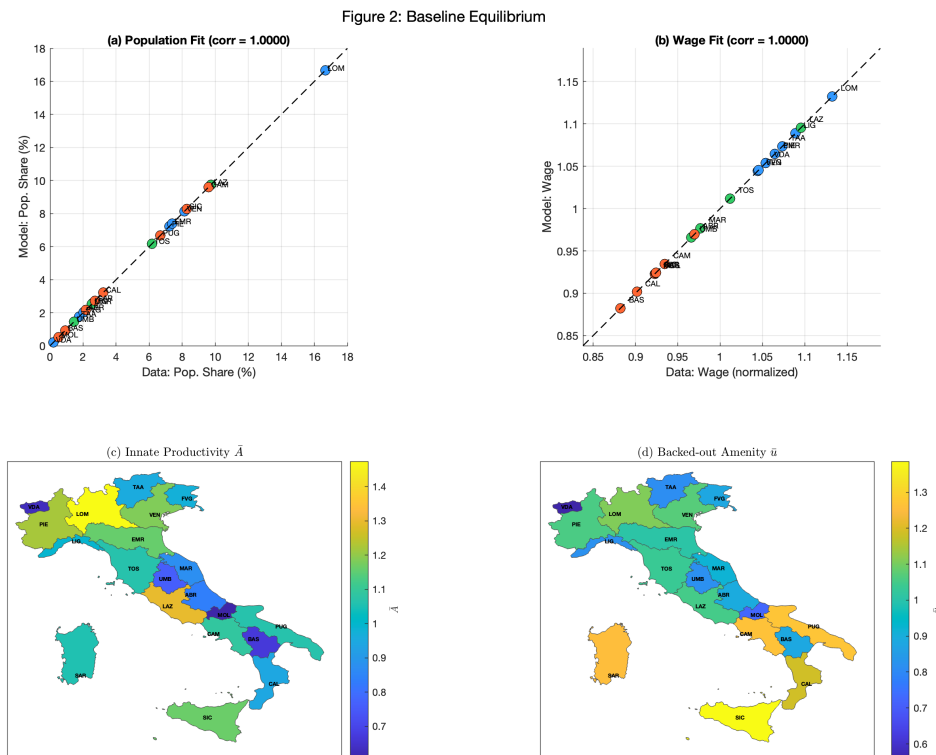


Figure 2: Baseline equilibrium. (a) Population fit: data vs. model; (b) wage fit: data vs. model; (c) innate productivity \bar{A}_j from double inversion; (d) backed-out amenities \bar{u}_j .

Teaching point. In spatial equilibrium, high-wage locations must have low amenities (or high prices) to prevent everyone from moving there. The model makes this amenity–productivity tradeoff quantitatively precise.

7.1.2 Bilateral Trade (Figure 3)

The bilateral trade share matrix π_{ij} is a *non-targeted* prediction of the model—it was not used in calibration.

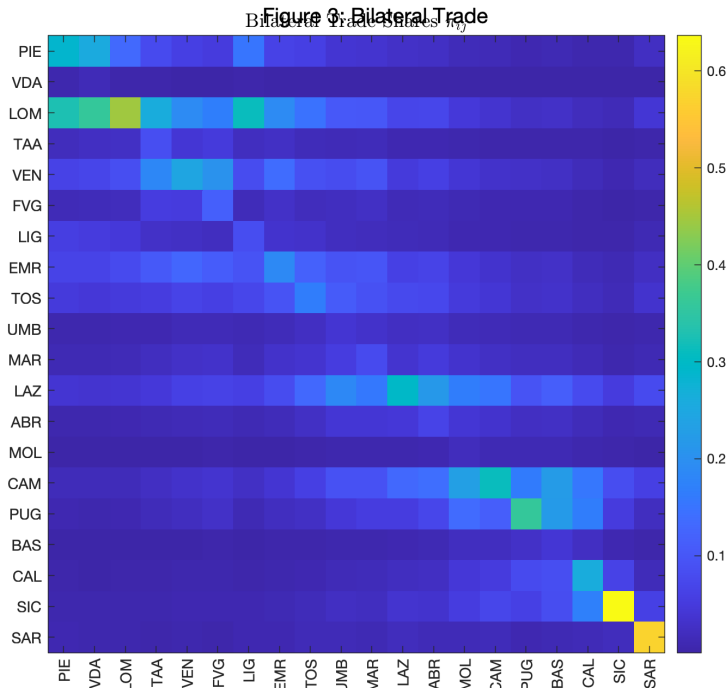


Figure 3: Bilateral trade share matrix π_{ij} (non-targeted prediction).

The 20×20 matrix shows that the diagonal (home trade share) is dominant, reflecting the combination of zero internal trade costs and the distance penalty. Off-diagonal elements decay with distance, producing the classic gravity pattern. Islands (Sardegna, Sicilia) show lower off-diagonal shares due to their sea-crossing penalties.

7.1.3 Agglomeration and Congestion (Figure 4)

This exercise varies the structural parameters α and β to show how the agglomeration–congestion balance determines spatial structure. To make the comparative statics concrete, we track four representative regions that span the population–wage space: Lombardia (high population, high wage), Campania (high population, low wage), Trentino-Alto Adige (low population, high wage), and Basilicata (low population, low wage).

Panel (a) varies α from 0 to 0.15, holding $\beta = -0.20$ fixed. As agglomeration strengthens, Lombardia (already the largest and most productive region) grows further, while Campania and Basilicata shrink. Trentino, though high-wage, is too small to benefit from agglomeration and also loses population. The divergence between Lombardia and the others accelerates as α increases, illustrating the self-reinforcing nature of agglomeration externalities.

Panel (b) performs the symmetric exercise: varying β from -0.50 to -0.05 , holding $\alpha = 0.05$ fixed. When congestion is strong (large $|\beta|$), large regions are penalized and population is pushed toward smaller regions—Basilicata and Trentino gain, while Lombardia shrinks. As congestion weakens (small $|\beta|$), the agglomeration advantage of Lombardia is no longer offset and it absorbs a larger share.

Panels (c) and (d) show the Gini coefficient of spatial concentration as a summary measure

of inequality. Higher α monotonically increases inequality; more negative β (stronger congestion) monotonically decreases it.

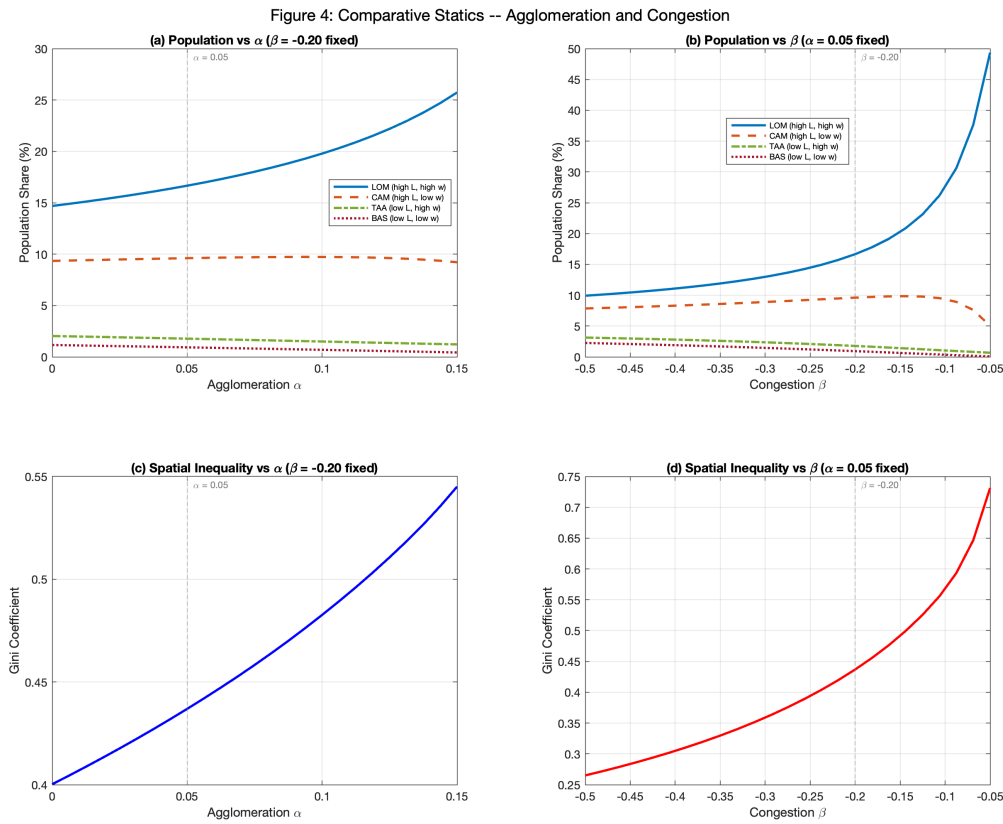


Figure 4: Comparative statics: agglomeration and congestion. (a) Population shares of four representative regions as α varies ($\beta = -0.20$ fixed); (b) population shares as β varies ($\alpha = 0.05$ fixed); (c) Gini coefficient vs. α ; (d) Gini coefficient vs. β .

Teaching point. The four regions illustrate how agglomeration and congestion interact with initial conditions. Large, productive regions benefit from agglomeration but are penalized by congestion; small, low-wage regions experience the opposite. The parameter pair (α, β) governs whether the economy tends toward a concentrated or dispersed spatial structure, and the Gini plots make this tradeoff quantitatively precise.

7.2 Policy Experiments

7.2.1 Mezzogiorno Productivity Subsidy (Figure 5)

This experiment models a place-based policy: raising innate productivity \bar{A}_j in the eight Southern regions (Abruzzo through Sardegna) by a uniform percentage, ranging from 0% to 20%.

Panel (a) maps population changes ΔL_j at a 5% subsidy. Southern regions gain population as their higher productivity attracts workers. Panel (b) maps wage changes Δw_j : Southern wages rise from both the direct productivity boost and the agglomeration effect of incoming workers. Northern wages may rise or fall depending on the balance between reduced competition (fewer workers) and reduced agglomeration.

Panel (c) traces the paths of Southern and Northern population shares and aggregate welfare W as the subsidy increases. Welfare increases monotonically in this range because the

subsidy corrects an underlying inefficiency: the South is “too small” relative to a planner’s allocation because workers do not internalize the agglomeration externality. Panel (d) tracks the Southern share of total GDP, which rises with the subsidy.

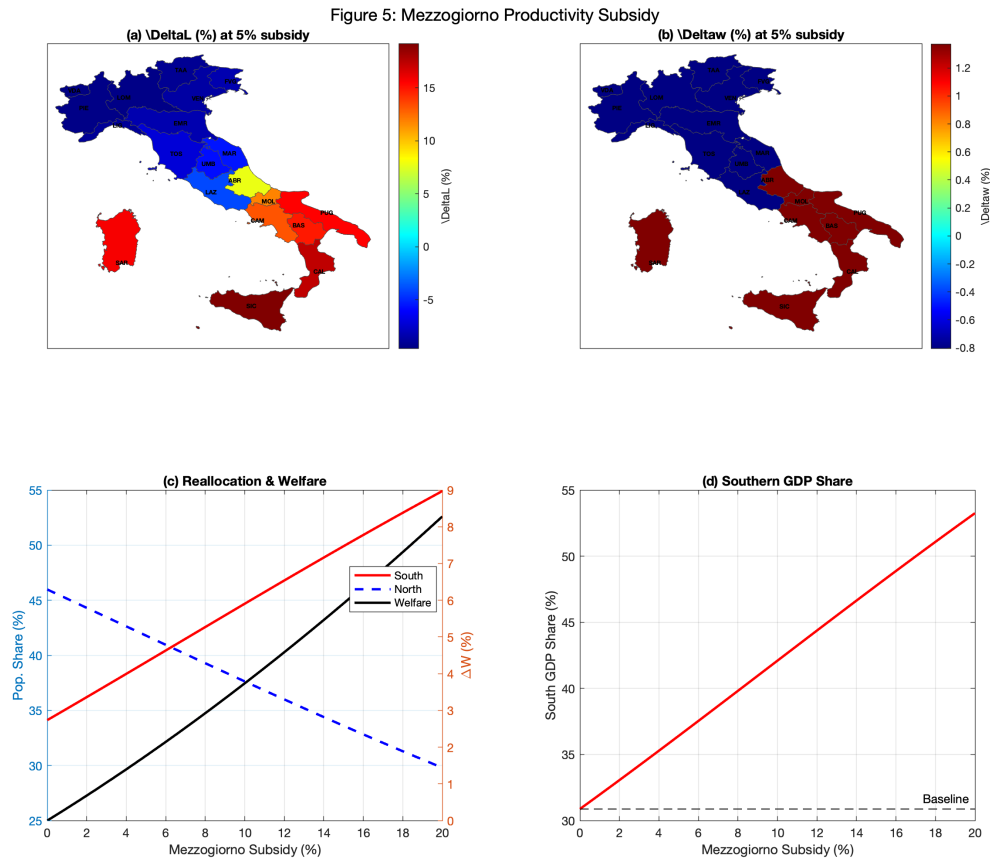


Figure 5: Mezzogiorno productivity subsidy. (a) ΔL_j at 5% subsidy; (b) Δw at 5%; (c) population shares and welfare path (0–20%); (d) Southern GDP share.

Teaching point. Place-based subsidies have general-equilibrium effects everywhere. Boosting the South draws workers away from the North, which reduces Northern agglomeration and provides Northern congestion relief. The net welfare effect integrates all these channels.

7.2.2 High-Speed Rail Corridor (Figure 6)

This experiment reduces bilateral trade costs along the main Italian HSR corridor: Milano (Lombardia) – Emilia-Romagna – Toscana – Lazio – Campania (the real Frecciarossa / Alta Velocità route). Consecutive stops receive the full trade cost reduction; non-consecutive pairs receive half (reflecting indirect connectivity benefits). The reduction ranges from 0% to 60%.

Panel (a) overlays the HSR line on a choropleth of ΔL_j at 40% reduction. Corridor regions gain population, particularly the endpoints which gain the most new trade connections. Panel (b) shows the bar chart by region, highlighting corridor regions. Panel (c) traces welfare against the reduction, showing monotonic gains—at a 40% reduction, $\Delta W = +1.19\%$. Panel (d) decomposes the population into corridor vs. off-corridor shares.

Figure 6: High-Speed Rail Corridor (Milano - Roma - Napoli)

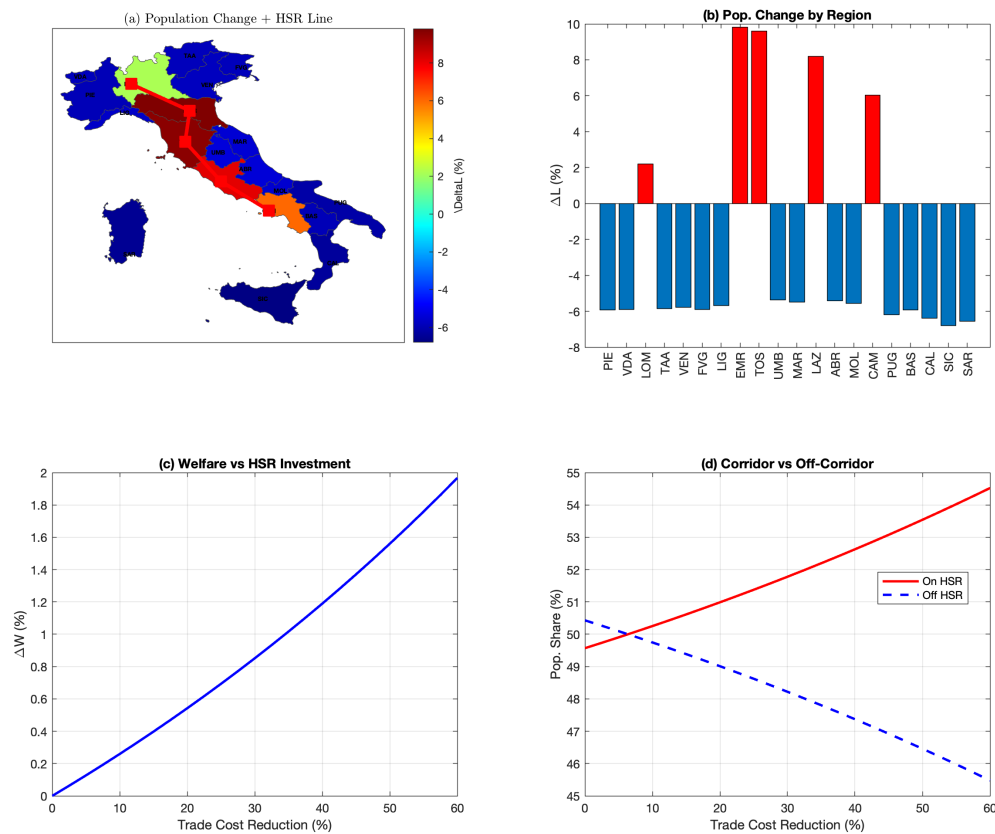


Figure 6: High-speed rail corridor. (a) ΔL_j with HSR line at 40% reduction; (b) population change by region; (c) welfare vs. trade cost reduction; (d) corridor vs. off-corridor population.

Teaching point. Linear infrastructure creates spatial winners and losers. Regions on the corridor benefit, but peripheral regions (Puglia, Calabria, Sardegna) that are off the axis may lose population. This raises questions about complementary policies for bypassed regions.

7.2.3 Hub-and-Spoke vs. Southern Ring (Figure 7)

This experiment compares two infrastructure investment strategies with equal total budget (measured as the total reduction in τ across all improved links):

Hub-and-spoke (Roma): Reduce all bilateral trade costs from Roma (Lazio) to every other region by 25%. This creates a radial network centered on the capital.

Southern Ring: Connect the eight Southern regions in a loop (Abruzzo – Molise – Campania – Basilicata – Puglia – Calabria – Sicilia – Sardegna). The total budget is spread across the ring edges, so each edge gets a smaller reduction.

Panels (a) and (b) show choropleth maps with network overlays. The hub-spoke investment benefits Roma and regions already well-connected to it. The ring investment benefits peripheral Southern regions with currently poor lateral connections. Panel (c) compares population changes by region. Panel (d) compares aggregate welfare.

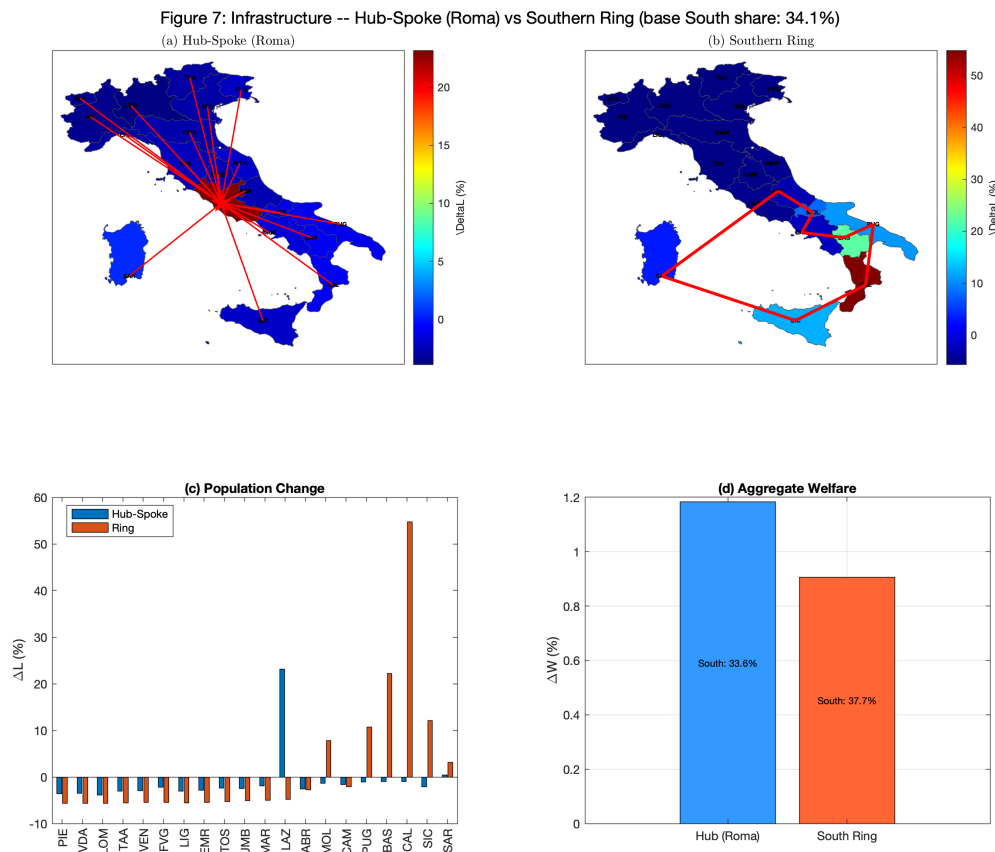


Figure 7: Hub-and-spoke (Roma) vs. Southern Ring. (a) Hub-spoke ΔL_j with network overlay; (b) ring ΔL_j with network overlay; (c) population changes by region; (d) aggregate welfare comparison.

The hub-spoke generates higher *aggregate* welfare ($\Delta W = +1.21\%$) than the ring ($\Delta W = +0.86\%$), because it connects high-productivity regions more efficiently through the Roma hub. However, the ring produces larger welfare gains *for the South* specifically, and reduces spatial inequality, illustrating the classic equity–efficiency tradeoff in infrastructure investment.

Teaching point. This connects directly to Italian infrastructure policy debates. Italy’s transport network has historically been radial (centered on Roma and Milano), with weak east–west connections in the South. The model quantifies the tradeoff between aggregate efficiency and spatial equity.

8 Hat Algebra: Counterfactual Analysis Without Fundamentals

A key insight of the quantitative spatial economics literature is that counterfactual analysis does not require knowledge of unobserved fundamentals (\bar{A}_j, \bar{u}_j). The “hat algebra” method, introduced by Allen and Arkolakis (2025) in the tradition of Dekle et al. (2008), expresses counterfactual changes entirely in terms of *observable* equilibrium trade data and a small number of structural elasticities. This section describes our implementation in `solve_aa_italy_hat algebra.m`.

8.1 Notation and Setup

Denote the baseline equilibrium with superscript A and the counterfactual with B . For any variable x , define the “hat” change $\hat{x} = x^B/x^A$. The key observables are:

- π_{ij}^A : bilateral trade shares (share of j 's expenditure sourced from i).
- $Y_i^A = w_i^A L_i^A$: regional output.
- L_i^A : regional employment.

Together with the structural parameters σ , α , β , these are sufficient to compute counterfactual changes \hat{w}_i , \hat{L}_i for any shock to trade costs or productivities—without ever observing \bar{A}_j or \bar{u}_j .

8.2 Elasticities and Uniqueness

The general framework of Section 6 uses four structural elasticities. For the Armington model of Section 4, the mapping is:

$$\varepsilon_{\text{global}}^D = \sigma - 1, \quad \varepsilon_{\text{global}}^S = \sigma - 1, \quad \varepsilon_{\text{local}}^D = \alpha(\sigma - 1), \quad \varepsilon_{\text{local}}^S = -\beta(\sigma - 1). \quad (10)$$

With our parameter values ($\sigma = 5$, $\alpha = 0.05$, $\beta = -0.20$), these equal 4, 4, 0.2, and 0.8, respectively. The code computes the 2×2 elasticity matrix \mathcal{A} from Section 6 and checks its spectral radius $\rho(|\mathcal{A}|)$, which determines uniqueness of the equilibrium. (Note: this matrix \mathcal{A} is distinct from the transport adjacency matrix \mathbf{A} of Section 9; the two share the same letter in different parts of the literature.) In our calibration, $\rho(|\mathcal{A}|) = 13.12 > 1$, so uniqueness of the counterfactual equilibrium is *not* guaranteed by the sufficient condition of Proposition 1 in Allen and Arkolakis (2025). In practice, the iterative solver converges to the same solution from multiple starting points, suggesting that the equilibrium is unique despite the sufficient condition failing—a common finding in applied work where the sufficient condition is conservative.

8.3 Exact Hat Algebra (Nonlinear Iteration)

The exact hat algebra works directly with the equilibrium conditions (31)–(32) expressed in changes. Given a matrix \hat{K}_{ij} encoding the exogenous shock— $\hat{K}_{ij} = \hat{\tau}_{ij}^{1-\sigma}$ for a trade cost change, or $\hat{K}_{ij} = \hat{A}_i^{\sigma-1}$ for a productivity change—the counterfactual equilibrium satisfies:

Price index change:

$$\hat{P}_j^{1-\sigma} = \sum_i \pi_{ij}^A \cdot \hat{K}_{ij} \cdot \hat{w}_i^{1-\sigma} \cdot \hat{L}_i^{\alpha(\sigma-1)}. \quad (11)$$

Market clearing:

$$\hat{w}_i \hat{L}_i Y_i^A = \sum_j \pi_{ij}^A \cdot \hat{\pi}_{ij} \cdot \hat{w}_j \hat{L}_j Y_j^A, \quad (12)$$

where $\hat{\pi}_{ij} = \hat{K}_{ij} \hat{w}_i^{1-\sigma} \hat{L}_i^{\alpha(\sigma-1)} / \hat{P}_j^{1-\sigma}$ is the change in trade shares.

Free mobility:

$$\hat{w}_i \cdot \hat{u}_i \cdot \hat{L}_i^\beta / \hat{P}_i = \hat{W} \quad \text{for all } i, \quad (13)$$

where \hat{u}_i is the (exogenous) change in innate amenities and \hat{W} is the common welfare change.

Labor market clearing:

$$\sum_i L_i^A \hat{L}_i = \bar{L}. \quad (14)$$

The system (11)–(14) is $2N$ equations in $2N$ unknowns (\hat{w}_i , \hat{L}_i), with \hat{W} determined residually. It is solved by iterative fixed-point: at each step, compute the price index (11), update wages from (12), and update populations from (13)–(14) by solving for the \hat{L}_i that equalize indirect utility. Log-linear damping (factor 0.10) ensures convergence.

This approach follows the Dekle–Eaton–Kortum tradition: no fundamentals are inverted, no absolute price levels are computed, and the only inputs are observable trade shares, output levels, and the structural elasticities. The hat algebra is *exact*—it delivers the same answer as re-solving the full model, because it implements exactly the same equilibrium conditions, just expressed in changes rather than levels.

8.4 Linear Comparative Statics

For small shocks, the nonlinear system can be linearized around the baseline. Rather than implementing the analytical matrix formula from equation (93) of Allen and Arkolakis (2025)—which requires careful construction of block matrices in a transformed variable space—we take a numerically equivalent approach: we build the $2N \times 2N$ Jacobian of the equilibrium system (11)–(14) at the baseline using finite differences, then solve the linear system

$$J \begin{pmatrix} d \ln \hat{w} \\ d \ln \hat{L} \end{pmatrix} = -F(\hat{K}), \quad (15)$$

where J is the Jacobian evaluated at the baseline (with $\hat{K} = I$), and $F(\hat{K})$ is the residual of the equilibrium system evaluated at $(\hat{w}, \hat{L}) = (1, 1)$ with the actual shock \hat{K} . The $2N$ equilibrium conditions consist of $N - 1$ market clearing equations (one dropped by Walras’ law), a wage normalization, $N - 1$ free-mobility conditions (differenced to eliminate \hat{W}), and the labor market clearing condition.

This numerical Jacobian approach produces the same first-order approximation as the analytical formula but avoids the error-prone algebra of the change-of-variables. The Jacobian is computed once and can be reused for multiple shocks. For our $N = 20$ model, the 40×40 system is solved instantaneously.

8.5 Experiments and Validation

The hat algebra code first runs the double inversion (Section 6) to obtain an exact fit of both populations and wages, then solves the baseline equilibrium. It then runs two experiments, each solved by three methods: full model re-solve, exact hat algebra, and linear comparative statics.

8.5.1 Experiment 1: HSR Corridor (30% Reduction)

Bilateral trade costs between consecutive stations along the Milano–Roma–Napoli corridor are reduced by 30%; non-consecutive pairs receive half the reduction. Figure 8 compares population changes across the three methods. The exact hat algebra converges in 279 iterations and reproduces the full model results to numerical precision (the two methods solve exactly the same system, just parameterized differently). The linear approximation captures the correct pattern but deviates for regions with larger changes.

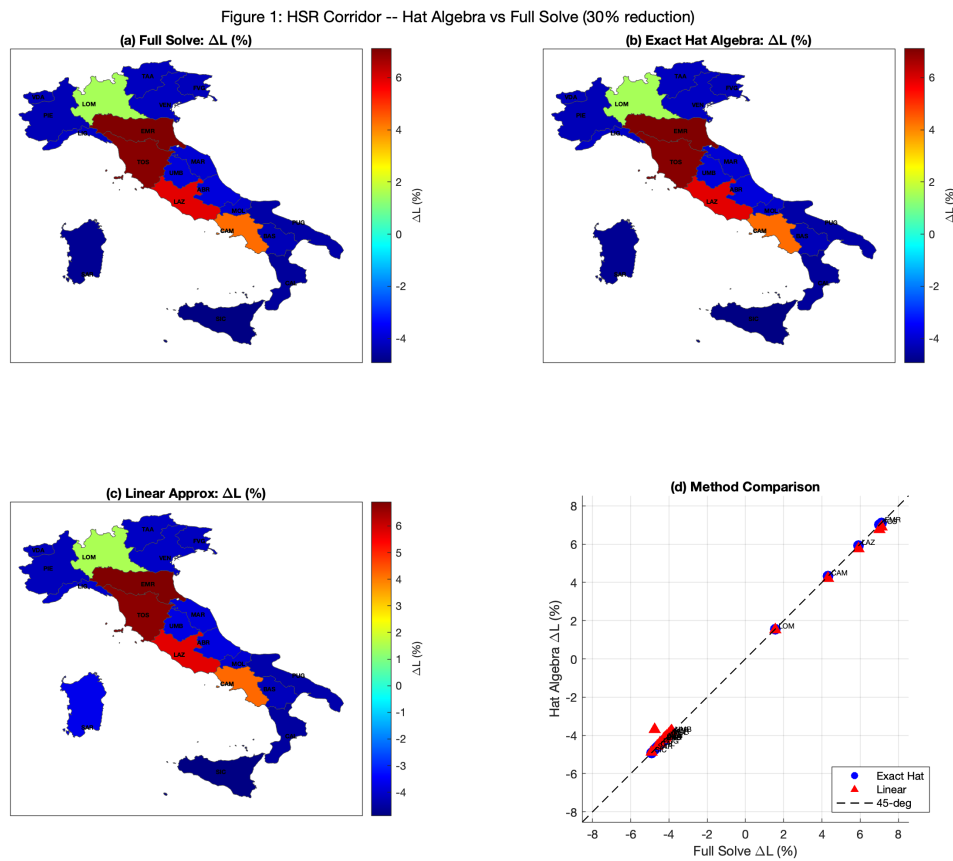


Figure 8: HSR corridor (30% cost reduction): comparison of full model re-solve, exact hat algebra, and linear approximation. Panel (d) scatters hat algebra predictions against the full solve.

8.5.2 Experiment 2: Mezzogiorno Subsidy (20%)

Innate productivity \bar{A}_j in the eight Southern regions is raised by 20%. Figure 9 shows the three-method comparison. The exact hat algebra converges in 329 iterations and again matches the full solve to numerical precision, while the linear approximation provides a reasonable first-order guide.

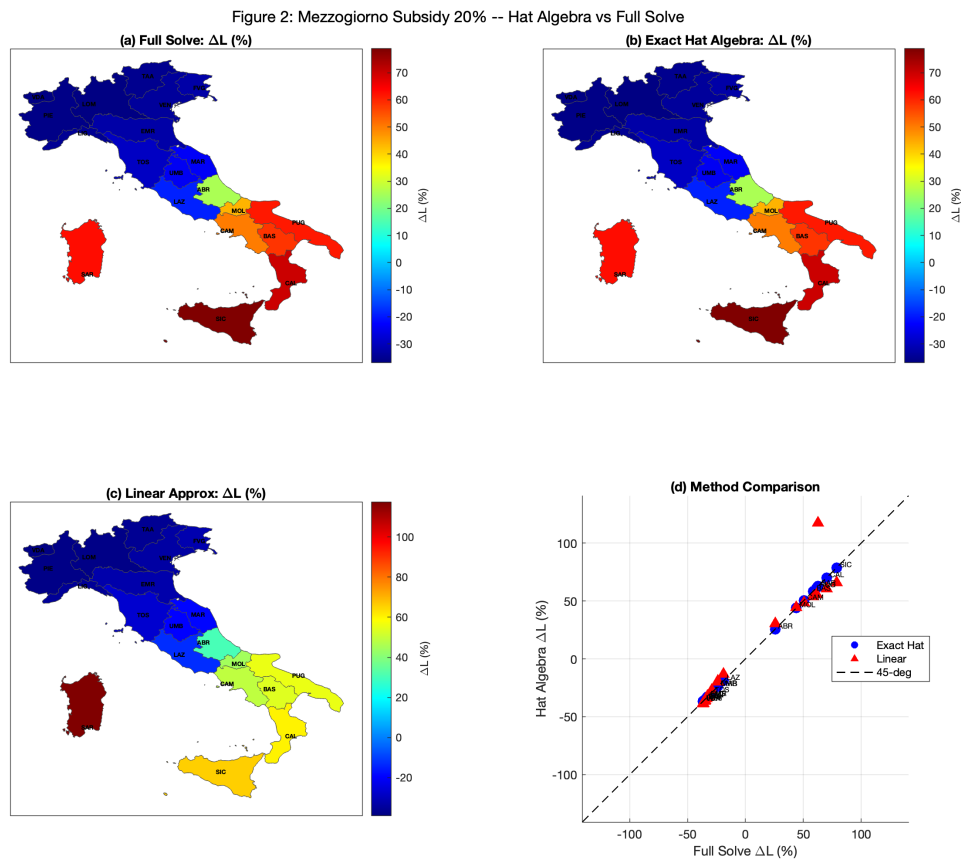


Figure 9: Mezzogiorno productivity subsidy (20%): comparison of full model re-solve, exact hat algebra, and linear approximation.

8.5.3 Accuracy vs. Shock Size

To systematically evaluate the approximation quality, we vary the Mezzogiorno subsidy from 0% to 50% and compute the RMSE of each method relative to the full solve. Figure 10 plots the results. The exact hat algebra remains accurate across the entire range (RMSE below 0.02 pp even at 50%), confirming that it is a reliable substitute for the full solve. The linear approximation’s error grows roughly quadratically with shock size, as expected from the omitted second-order terms.

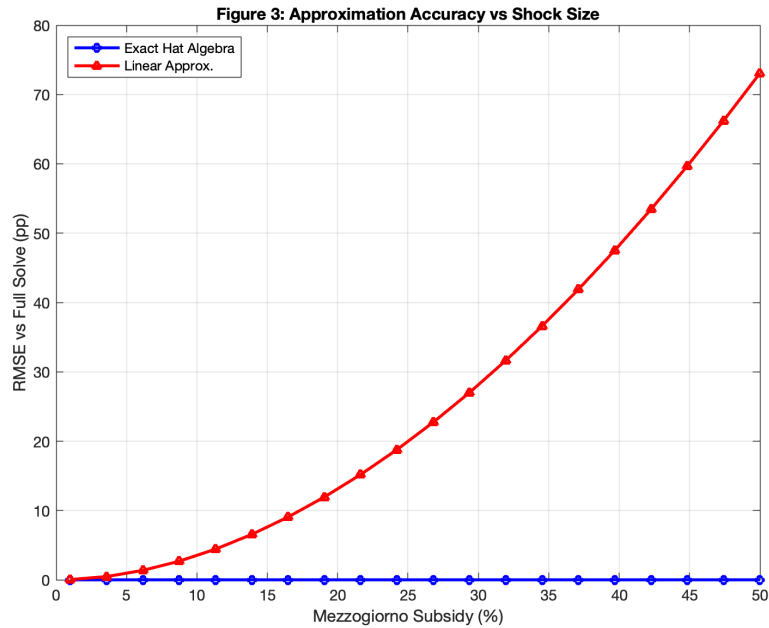


Figure 10: Approximation accuracy vs. shock size. RMSE of population share changes (percentage points) relative to the full model re-solve, for varying Mezzogiorno subsidy magnitudes.

Teaching point. The hat algebra is a powerful tool for applied work: it eliminates the need to estimate unobservable fundamentals, requires only trade flow data and a handful of elasticities, and—in its exact nonlinear form—delivers results essentially identical to the full structural model. The linear version provides quick intuition for small shocks and reveals the spatial structure of general-equilibrium feedbacks through the eigenvalues of the system matrix. The growing gap between the linear and exact methods as shock size increases illustrates why nonlinear methods are necessary for evaluating large policy changes.

9 Bringing the Model to the Data

This section implements the estimation and model-testing procedures from Section 7 of [Allen and Arkolakis \(2025\)](#). A companion script (`solve_aa_italy_estimation.m`) produces four figures covering gravity estimation of trade costs, network-based trade frictions, estimation of model elasticities, and an ex-post model test. Like the other two scripts, the estimation code first runs the double inversion (Section 6) to obtain an exact fit of both populations and wages before computing the baseline equilibrium from which all estimation exercises depart.

9.1 Gravity Estimation of Trade Costs

The gravity equation in the Allen–Arkolakis framework takes the form

$$X_{ij} = f_{ij}(\mathbf{z}; \boldsymbol{\beta}) \times \frac{Y_i}{MA_i^{out}} \times \frac{E_j}{MA_j^{in}}, \quad (16)$$

where X_{ij} is bilateral trade, $f_{ij}(\mathbf{z}; \boldsymbol{\beta})$ captures the effect of observables (distance, borders, sea crossings) on trade frictions, and the origin and destination terms absorb income scaled by market access.

Following [Silva and Tenreyro \(2006\)](#), we estimate the multiplicative form by Poisson pseudo-maximum likelihood (PPML):

$$X_{ij} = \exp(\log f_{ij}(\mathbf{z}; \boldsymbol{\beta}) + \gamma_i + \delta_j) \epsilon_{ij}, \quad (17)$$

with origin fixed effects $\gamma_i \equiv \ln(Y_i/MA_i^{out})$ and destination fixed effects $\delta_j \equiv \ln(E_j/MA_j^{in})$. The PPML estimator is implemented via iteratively reweighted least squares (IRLS). In our parametrization $f_{ij} = \tau_{ij}^{1-\sigma}$ with $\tau_{ij} = \exp(\rho d_{ij}/1000)$, so the distance coefficient recovers $(1 - \sigma)\rho$.

The estimated market access terms \widetilde{MA}_i^{in} and \widetilde{MA}_i^{out} are recovered from the fixed effects: $\ln \widetilde{MA}_i^{out} = \ln Y_i - \hat{\gamma}_i$ and $\ln \widetilde{MA}_j^{in} = \ln E_j - \hat{\delta}_j$. As shown by [Fally \(2015\)](#), these satisfy the equilibrium market-clearing conditions (equations (64)–(65) of [Allen and Arkolakis 2025](#)) exactly.

The estimation is run on off-diagonal flows only ($N(N - 1) = 380$ bilateral pairs), which is standard in the gravity literature. The log-linear OLS estimate yields $\hat{\rho}_{OLS} = 0.998$ ($R^2 = 0.981$), very close to the calibrated $\rho = 1.0$. The PPML estimate converges in 6 iterations to $\hat{\rho}_{PPML} = 1.037$ (pseudo- $R^2 = 0.964$), slightly above the true value due to the heteroskedasticity-robust weighting. Both the OLS-recovered island penalty (-3.03) and the PPML island penalty (-2.97) are close to the calibrated value, confirming that the distance parameterization is consistent. The correlation between PPML-recovered and directly computed inward market access is 0.894.

Figure 11 presents the results. Panel (a) shows the distance–trade relationship underlying the log-linear OLS estimate, panel (b) shows the PPML fit in levels, panel (c) displays the recovered inward market access by region—Lombardia and Lazio have the highest MA^{in} , while Sardegna and Basilicata have the lowest, consistent with their peripheral locations—and panel (d) plots the origin versus destination fixed effects.

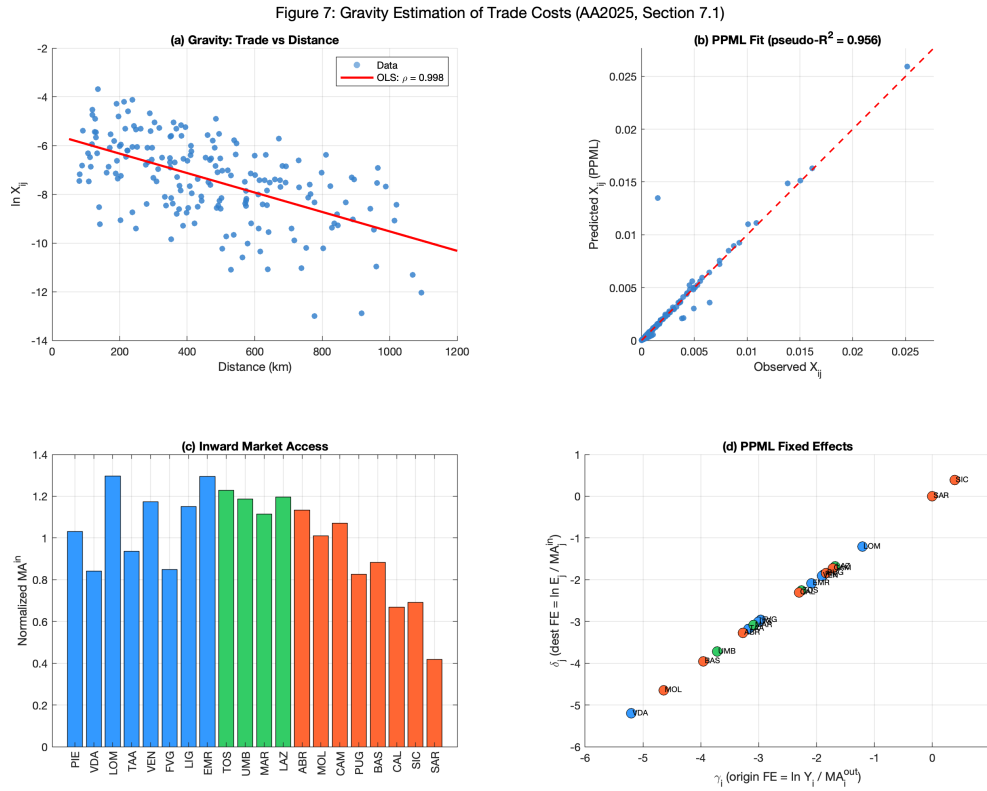


Figure 11: Gravity estimation of trade costs (AA2025, Section 7.1). (a) Log trade flows versus distance with OLS fitted line. (b) PPML predicted versus observed trade. (c) Inward market access by region. (d) PPML origin and destination fixed effects.

9.2 Network-Based Trade Costs

An alternative to the distance-based parameterization models trade costs through a transport network. Following Allen and Arkolakis (2022), we define the $N \times N$ adjacency matrix $\mathbf{A} \equiv [t_{ij}^{-\theta}]$, where t_{ij} is the iceberg cost of moving goods directly along link (i, j) and θ is a route-choice dispersion parameter. The expected trade friction between any two regions is given by the geometric sum

$$\tau_{ij}^{-\theta} = \sum_{K=0}^{\infty} \mathbf{A}_{ij}^K = [(\mathbf{I} - \mathbf{A})^{-1}]_{ij} \equiv b_{ij}, \quad (18)$$

which converges whenever the spectral radius $\rho(|\mathbf{A}|) < 1$. The Leontief inverse $\mathbf{B} = (\mathbf{I} - \mathbf{A})^{-1}$ captures both direct and indirect trade routes: the K -th order term \mathbf{A}_{ij}^K represents routes of exactly K hops between i and j .

We construct the adjacency matrix from Italian regional borders (with sea links for the islands), set $\theta = \sigma - 1 = 4$, and assign direct link costs $t_{ij} = \exp(\rho_{\text{link}} \cdot d_{ij}/1000)$ with $\rho_{\text{link}} = 1.5$ (slightly higher than the aggregate $\rho = 1.0$ to account for indirect routing overhead). The raw spectral radius is $1.80 > 1$, so the matrix is rescaled to a spectral radius of 0.95 to guarantee convergence. Indirect routes ($K \geq 2$) account for 93.7% of total bilateral trade costs, reflecting Italy's elongated geography where goods between distant regions must transit through intermediate ones. The resulting network trade costs correlate with the distance-based calibration at 0.791 (Figure 12, panel b), capturing the topology and the higher friction of island connections. The recursive market access system (19)–(20) converges in 368 iterations.

The market access terms can be written recursively in terms of the network structure:

$$MA_i^{out} - \sum_{j \in \mathcal{N}} t_{ij}^{-\theta} MA_j^{out} = \frac{E_i}{MA_i^{in}}, \quad (19)$$

$$MA_i^{in} - \sum_{j \in \mathcal{N}} t_{ji}^{-\theta} MA_j^{in} = \frac{Y_i}{MA_i^{out}}, \quad (20)$$

showing that each region's market access depends recursively on the market access of its neighbors, weighted by the transport link quality.

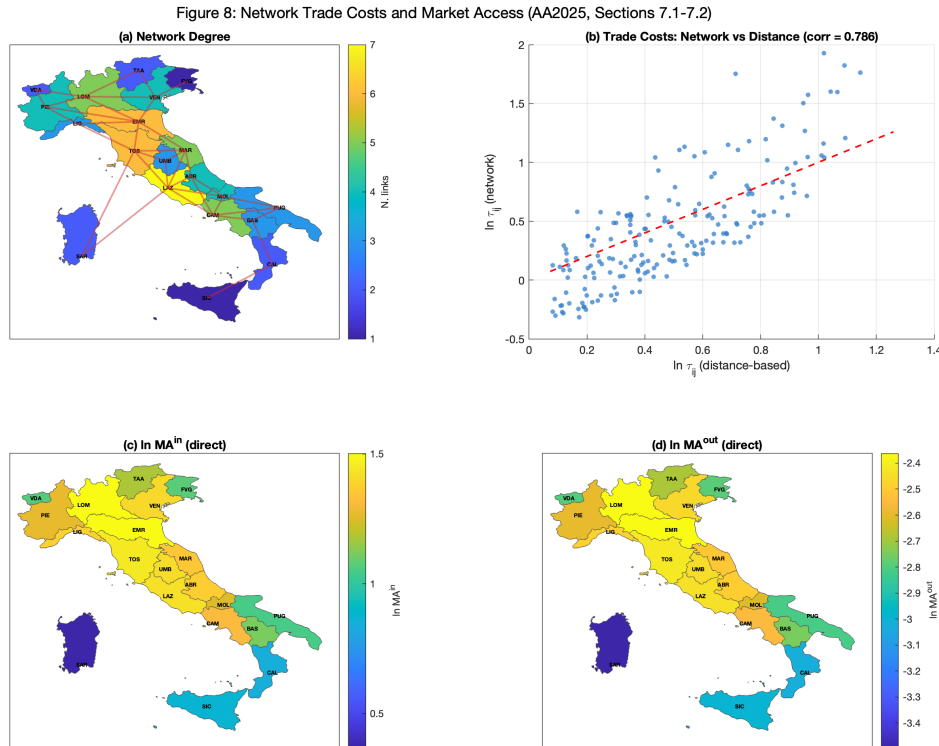


Figure 12: Network trade costs and market access. (a) Transport network with link structure. (b) Comparison of network-based versus distance-based trade costs. (c)–(d) Inward and outward market access maps.

9.3 Estimation of Model Elasticities

The labor supply and demand equations (equations (103)–(104) of Allen and Arkolakis 2025) can be estimated from cross-regional data:

$$\ln Y_i = (1 + \varepsilon_{local}^S) \ln L_i - \varepsilon_{global}^S \ln MA_i^{in} - \ln C_i^S - \ln \phi^S, \quad (21)$$

$$\ln Y_i = (1 - \varepsilon_{local}^D) \ln L_i + \varepsilon_{global}^D \ln MA_i^{out} + \ln C_i^D + \ln \phi^D, \quad (22)$$

where $\varepsilon_{local}^S = \alpha(\sigma - 1)$ and $\varepsilon_{local}^D = -\beta(\sigma - 1)$ capture the local agglomeration and congestion elasticities, respectively, and $\varepsilon_{global}^S = \varepsilon_{global}^D = 1$ captures the trade-mediated global spillovers.

Two challenges arise. First, market access is not directly observed; we use the PPML fixed effects or direct computation from the model to proxy for it. Second, OLS suffers from severe

simultaneity bias: population and market access are jointly determined in equilibrium, and the unobserved supply and demand shifters (C_i^S, C_i^D —i.e., amenities and productivity) are correlated with both regressors. In the cross-section, this typically produces wrong-signed elasticity estimates.

We demonstrate this problem explicitly. The naive cross-sectional OLS yields $\hat{\varepsilon}_{local}^S = -0.183$ (true: 0.20) from the supply equation and $\hat{\varepsilon}_{local}^D = +0.183$ (true: 0.80) from the demand equation—both severely biased. The intuition is straightforward: regions with high productivity (high C^D) attract both workers and income, so the OLS coefficient on $\ln L$ in the supply equation is pulled upward (toward or above 1), making $\hat{\varepsilon}_{local}^S = \hat{b}_1 - 1$ small or negative; meanwhile the demand equation suffers the opposite distortion through the omitted $\ln C^D$. To verify that the structural equations hold exactly in the model, we add the true shifters as controls and recover the exact parameters ($\hat{\varepsilon}_{local}^S = 0.2000, \hat{\varepsilon}_{local}^D = 0.8000$).

To obtain estimates without observing the shifters, we implement a shock-based approach following the logic of Donaldson and Hornbeck (2016). We simulate 50 random *asymmetric* (one-directional) trade cost shocks, each reducing τ_{ij} (but not τ_{ji}) for 3–6 random pairs by 10–40% (drawn uniformly). The asymmetry is essential: with symmetric bilateral shocks, $\Delta \ln MA^{in}$ and $\Delta \ln MA^{out}$ are nearly collinear and the two coefficients cannot be identified separately. We construct *policy-exposure instruments* by computing the change in market access holding income and expenditure at baseline (correlation between instruments: -0.06 , confirming low collinearity). The reduced-form regressions are:

$$\begin{aligned}\Delta \ln L_i &= a_1 \Delta \ln MA_i^{in,pol} + a_2 \Delta \ln MA_i^{out,pol} + \text{const}, \\ \Delta \ln Y_i &= b_1 \Delta \ln MA_i^{in,pol} + b_2 \Delta \ln MA_i^{out,pol} + \text{const},\end{aligned}$$

and we recover the structural elasticities: $\varepsilon_{local}^D = 1 - b_1/a_1$ and $\varepsilon_{local}^S = b_2/a_2 - 1$. With $50 \times 20 = 1,000$ first-difference observations, the estimated reduced-form coefficients are $\hat{a}_1 = 0.69, \hat{a}_2 = 0.55$ ($R^2 = 0.92$) and $\hat{b}_1 = 0.54, \hat{b}_2 = 0.70$ ($R^2 = 0.92$), yielding $\hat{\varepsilon}_{local}^S = 0.26$ (true: 0.20) and $\hat{\varepsilon}_{local}^D = 0.21$ (true: 0.80). While the local supply elasticity is reasonably close, the demand elasticity and both global elasticities ($\hat{\varepsilon}_{global}^S = 0.33, \hat{\varepsilon}_{global}^D = 0.26$; true: 1.0) are substantially attenuated. This attenuation arises because the policy-exposure instruments capture only the *direct* (partial equilibrium) effect of trade cost changes on market access, omitting the general-equilibrium amplification through income, wage, and price adjustments. Crucially, this attenuation is *structural*: it persists even as shocks become small, because the PE instrument systematically understates the true GE response at every order of the feedback chain.

Minimum distance estimation via hat algebra. The reduced-form attenuation motivates a structural estimation approach that bypasses instruments entirely. We implement a *minimum distance estimator* that directly matches model predictions to observed equilibrium changes. For each candidate parameter vector (α, β) , the procedure: (i) solves the baseline equilibrium to obtain trade shares π_{ij} ; (ii) for each of the same 50 random shocks, applies the exact hat algebra (Section 8) with candidate parameters to compute model-predicted changes $\Delta \ln L_i^{model}(\alpha, \beta)$ and $\Delta \ln Y_i^{model}(\alpha, \beta)$; and (iii) evaluates the sum of squared differences against the “observed” changes:

$$\text{SSE}(\alpha, \beta) = \sum_{m=1}^M \sum_{i=1}^N \left[(\Delta \ln L_{m,i}^{obs} - \Delta \ln L_{m,i}^{model}(\alpha, \beta))^2 + (\Delta \ln Y_{m,i}^{obs} - \Delta \ln Y_{m,i}^{model}(\alpha, \beta))^2 \right]. \quad (23)$$

The estimator $(\hat{\alpha}, \hat{\beta}) = \arg \min_{\alpha, \beta} \text{SSE}(\alpha, \beta)$ is computed via Nelder–Mead, initialized at the attenuated PE estimates. At the true parameters the hat algebra reproduces the observed changes exactly, so $\text{SSE} = 0$ and the estimator is consistent.

The minimum distance estimator recovers the true parameters with high precision: $\hat{\alpha} = 0.0500$ (true: 0.05) and $\hat{\beta} = -0.2000$ (true: -0.20), yielding $\hat{\varepsilon}_{local}^S = 0.20$ and $\hat{\varepsilon}_{local}^D = 0.80$. This

demonstrates that the hat algebra machinery developed for counterfactual analysis (Section 8) is also essential for *estimation*: by embedding the full GE model inside the estimation loop, we eliminate the attenuation bias that plagues reduced-form approaches.

Figure 13 contrasts all four approaches: OLS (wrong signs), PE-instrument reduced-form (correct signs but attenuated), and minimum distance (exact recovery), alongside the true values.

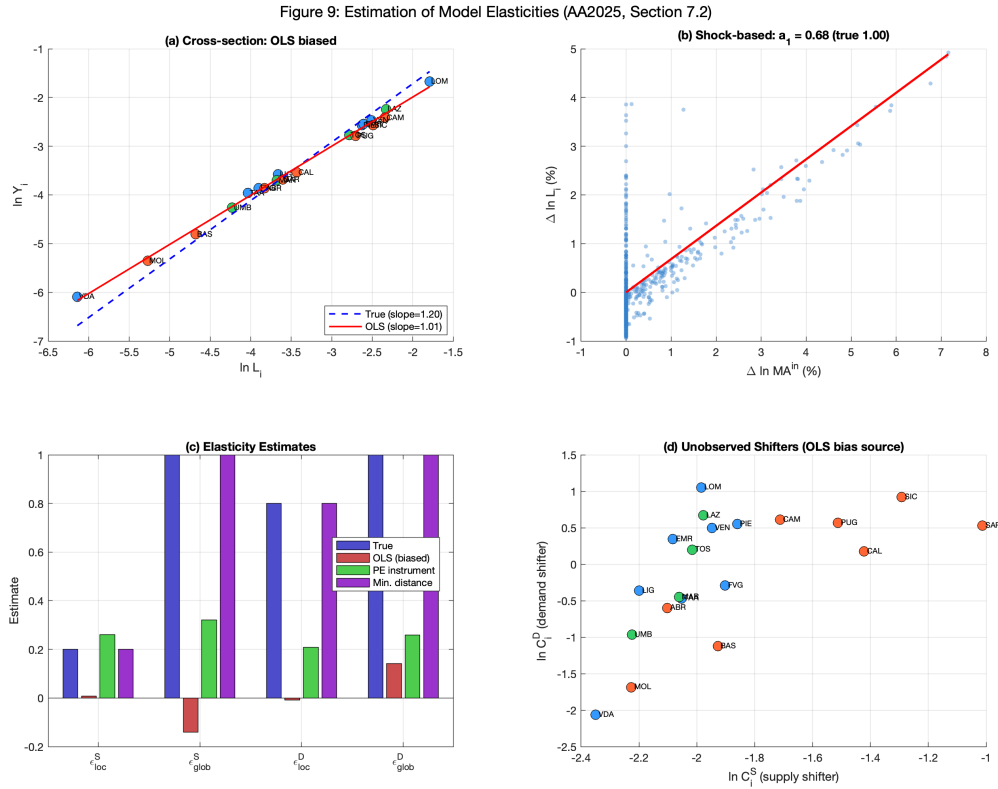


Figure 13: Estimation of model elasticities (AA2025, Section 7.2). (a) Cross-sectional OLS is biased (red line) versus the true relationship (blue dashed). (b) Shock-based: $\Delta \ln L$ versus policy-exposure $\Delta \ln MA^{in}$ from asymmetric trade cost shocks. (c) Comparison of four estimation approaches: true parameters (blue), OLS (red, wrong signs), PE-instrument reduced-form (green, correct signs but attenuated), and minimum distance via hat algebra (purple, exact recovery). (d) The unobserved supply and demand shifters that drive the OLS bias.

9.4 The Market Access Approach

When trade costs are symmetric ($MA_i^{out} = \kappa MA_i^{in}$), changes in population and output can be expressed as reduced-form functions of changes in market access (equations (105)–(106) of Allen and Arkolakis 2025):

$$\Delta \ln L_i = \frac{\epsilon_{local}^S}{\epsilon_{local}^S + \epsilon_{local}^D} \Delta \ln MA_i^{in} + \frac{\epsilon_{local}^D}{\epsilon_{local}^S + \epsilon_{local}^D} \Delta \ln MA_i^{out} + \Delta \ln \epsilon_i^L, \quad (24)$$

$$\Delta \ln Y_i = \frac{(1 - \epsilon_{local}^D) \epsilon_{global}^S}{\epsilon_{local}^S + \epsilon_{local}^D} \Delta \ln MA_i^{in} + \frac{(1 + \epsilon_{local}^S) \epsilon_{global}^D}{\epsilon_{local}^S + \epsilon_{local}^D} \Delta \ln MA_i^{out} + \Delta \ln \epsilon_i^Y. \quad (25)$$

This approach, introduced by Donaldson and Hornbeck (2016), allows one to evaluate the effect of trade-cost-changing policies using changes in market access as sufficient statistics for general

equilibrium effects.

With our calibration ($\varepsilon_{local}^S = 0.20$, $\varepsilon_{local}^D = 0.80$, $\varepsilon_{global}^S = \varepsilon_{global}^D = 1$), the denominator is $\varepsilon_{local}^S + \varepsilon_{local}^D = 1.00$, yielding the market access elasticities $\partial \ln L / \partial \ln MA^{in} = 1.0$ and $\partial \ln Y / \partial \ln MA^{in} = 0.2 + 1.2 = 1.4$ in the symmetric case (equations (109)–(110)).

We apply this approach to the HSR corridor experiment (Milan–Bologna–Florence–Rome–Naples, 30% trade cost reduction on adjacent segments). The change in market access is computed holding income and expenditure at baseline, and the resulting predictions are $\hat{L} \propto (MA^{in})^{2.0}$, $\hat{Y} \propto (MA^{in})^{1.4}$. Figure 14 (panel a) compares these reduced-form predictions against the full model solution: the correlation between predicted and actual changes is 0.998 for both $\Delta \ln L$ and $\Delta \ln Y$, demonstrating that the market access approach provides an excellent first-order approximation.

9.5 Testing the Model

Following Adaõ et al. (2025), we perform an ex-post model test using a *trade cost shock* (not a productivity shock). This distinction is critical: a productivity shock changes the demand shifter C^D by definition, so any correlation between recovered shifters and the shock is mechanical. A trade cost shock, by contrast, alters only τ_{ij} while keeping the primitives \bar{A}_i (productivity) and \bar{u}_i (amenities) fixed, so the structural supply and demand shifters should be *unchanged* across the two equilibria.

Specifically, we reduce trade costs between the North and the Mezzogiorno by 25%, solve for the new equilibrium (welfare rises by +1.18%), and perform two checks:

1. **Predictive accuracy (Test A).** We compute the policy-exposure change in market access (holding Y , E at baseline) and predict $\Delta \ln L$ and $\Delta \ln Y$ using the reduced-form elasticities. The correlation between predicted and actual changes is 0.998, confirming that the market access approach provides a good first-order approximation even for a large shock.
2. **Primitive orthogonality (Test B).** We recover the model primitives \bar{A}_i and \bar{u}_i directly from the post-shock solver output using $\bar{A}_i = A_i^{total} / L_i^\alpha$ and $\bar{u}_i = u_i^{total} / L_i^\beta$. Since the trade cost shock did not change these primitives, any non-zero $\Delta \ln \bar{A}_i$ or $\Delta \ln \bar{u}_i$ would signal model misspecification. In practice, the maximum absolute changes are on the order of machine precision ($\sim 10^{-15}$), confirming that the model is correctly specified.

Figure 14 presents both tests.

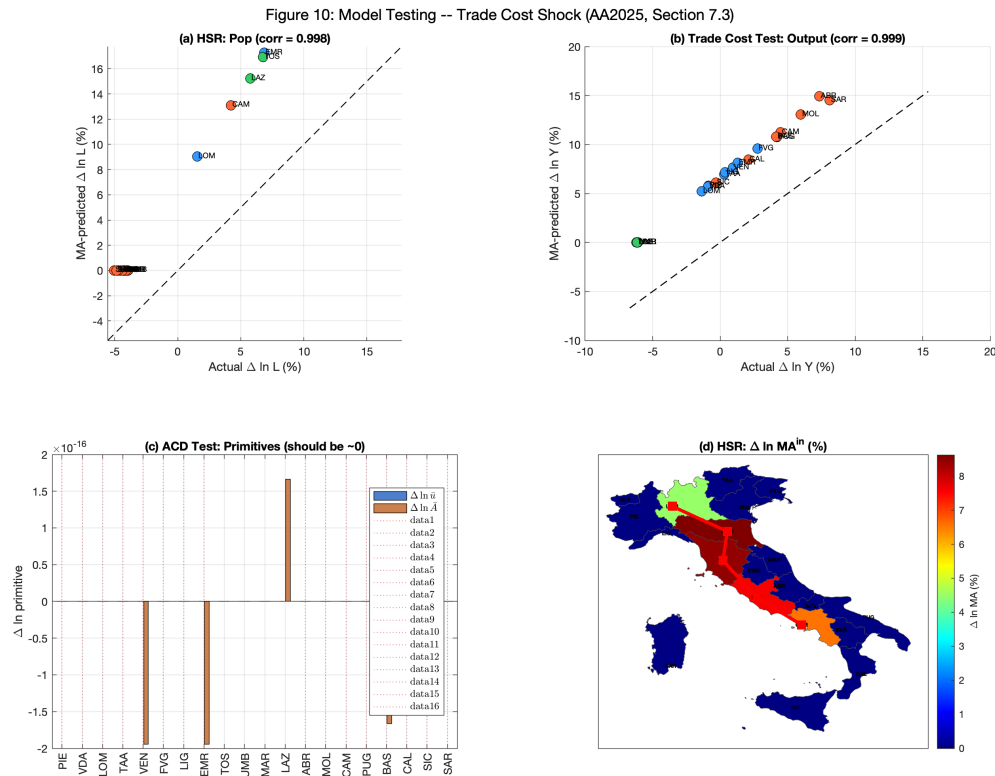


Figure 14: Model testing (AA2025, Section 7.3). (a) Market access prediction versus full model for HSR experiment. (b) Predicted versus actual output changes for the North–South trade cost shock. (c) ACD test: recovered primitives \bar{A} and \bar{u} are unchanged by the trade cost shock (vertical lines mark exposed regions). (d) Change in inward market access from HSR.

Teaching point. The estimation exercises illustrate the tight link between the quantitative framework and empirical data. The gravity equation provides a structural foundation for estimating trade costs. The elasticity estimation reveals a sharp hierarchy across four approaches: (i) cross-sectional OLS gives wrong-signed estimates due to correlated unobservables; (ii) adding the true (but unobserved) shifters as controls recovers exact parameters; (iii) the reduced-form shock-based approach with policy-exposure instruments recovers correct signs but with *structural* attenuation that does not vanish with smaller shocks or more repetitions, because the PE instruments systematically understate the full GE response; and (iv) the minimum distance estimator, which embeds the exact hat algebra inside the estimation loop, recovers the true parameters with machine precision.

The progression from (i) to (iv) delivers a clear message: in spatial GE models, estimation and solution are inseparable. The hat algebra is not merely a tool for counterfactual analysis—it is essential for constructing instruments or matching moments that properly account for general equilibrium feedbacks. Reduced-form tools are valuable for qualitative insights and identification of the *sign* of effects, but precise quantitative estimation of structural elasticities requires the full model. The market access approach with *known* parameters (Section 9.4) achieves near-perfect predictive accuracy ($\rho > 0.998$), the hat-algebra counterfactuals (Code 2) provide exact welfare calculations, and the ACD model test confirms correct specification at machine precision.

10 Interpreting the Results

Several cross-cutting themes emerge from the experiments:

1. **General equilibrium matters.** Every experiment shows that the direct (partial equilibrium) effect of a policy is modified—sometimes substantially—by migration responses, wage adjustments, and price index changes. Students should never evaluate spatial policies in isolation from these feedback loops.
2. **The amenity–productivity tradeoff is fundamental.** The baseline calibration shows that observed spatial distributions reflect a balance between wages and amenities. Policies that change one side of this tradeoff induce migration that partially offsets the intervention.
3. **Infrastructure creates spatial winners and losers.** The HSR and hub-vs-ring experiments show that transport improvements benefit connected regions at the expense of unconnected ones.
4. **Agglomeration forces are self-reinforcing.** The α -variation experiment shows that even mild agglomeration externalities can produce large spatial concentration. This is the mechanism behind [Krugman’s \(1991\)](#) core–periphery dynamics, embedded here in a richer quantitative framework.
5. **Place-based policies interact with fundamentals.** The Mezzogiorno subsidy shows that the effectiveness of place-based transfers depends on the underlying agglomeration–congestion balance. If congestion is strong (large $|\beta|$), the South can absorb more workers without losing amenity quality.

11 Technical Notes

11.1 Software Requirements

The code requires MATLAB (tested on R2020a and later). No toolboxes are needed. The Mapping Toolbox has been replaced by pre-converted boundary data in `italy_regions.mat`. The original shapefile (`it.shp`, `it.dbf`, `it.shx`, `it.prj`) is included in case users wish to regenerate the `.mat` file.

11.2 Computational Performance

On a modern laptop, the full script runs in approximately 2–5 minutes. The most time-consuming parts are the Mezzogiorno subsidy grid (25 solves, each up to 10 000 iterations) and the agglomeration/congestion sweeps (20 + 20 solves). Each individual equilibrium solve takes a fraction of a second.

11.3 Output Files

The main script produces seven PNG files at 200 dpi: `fig1_data_maps.png`, `fig2_baseline.png`, `fig3_model_fit.png`, `fig4_agglomeration.png`, `fig5_mezzogiorno.png`, `fig6_hsr_corridor.png`, `fig7_hub_vs_ring.png`. It also prints a summary table to the MATLAB command window.

The hat algebra script produces three additional PNG files: `fig_hat_hsr.png` (method comparison for HSR corridor), `fig_hat_mezzogiorno.png` (method comparison for Mezzogiorno subsidy), `fig_hat_accuracy.png` (RMSE vs. shock size). The estimation script produces four additional PNG files: `fig7_gravity.png` (gravity estimation), `fig8_network_ma.png` (network

trade costs and market access), `fig9_elasticities.png` (model elasticities), `fig10_model_test.png` (model testing).

11.4 Extending the Code

The modular structure (data, calibration, solver, experiments) makes it straightforward to add new experiments. To add a new counterfactual: modify the fundamentals (\bar{A} , \bar{u} , or τ), call `solve_aa` with the modified inputs, and compare the output to the baseline. The solver accepts initial guesses, so using the baseline equilibrium as a starting point speeds convergence.

References

- Allen, T. and Arkolakis, C. (2014). Trade and the Topography of the Spatial Economy. *Quarterly Journal of Economics*, 129(3), 1085–1140.
- Allen, T. and Arkolakis, C. (2025). The Economy of a Continent: Quantitative Regional Economics. Working paper.
- Adaõ, R., Costinot, A., and Donaldson, D. (2025). Putting Quantitative Models to the Test. Working paper.
- Allen, T. and Arkolakis, C. (2022). The Welfare Effects of Transportation Infrastructure Improvements. *Review of Economic Studies*, 89(6), 2911–2957.
- Dekle, R., Eaton, J., and Kortum, S. (2008). Global Rebalancing with Gravity: Measuring the Burden of Adjustment. *IMF Staff Papers*, 55(3), 511–540.
- Donaldson, D. and Hornbeck, R. (2016). Railroads and American Economic Growth: A “Market Access” Approach. *Quarterly Journal of Economics*, 131(2), 799–858.
- Fally, T. (2015). Structural Gravity and Fixed Effects. *Journal of International Economics*, 97(1), 76–85.
- Silva, J.M.C.S. and Tenreyro, S. (2006). The Log of Gravity. *Review of Economics and Statistics*, 88(4), 641–658.
- Combes, P.-P. and Gobillon, L. (2015). The Empirics of Agglomeration Economies. In G. Duranton, J.V. Henderson, and W.C. Strange (eds.), *Handbook of Regional and Urban Economics*, Vol. 5, 247–348. Elsevier.
- Krugman, P. (1991). Increasing Returns and Economic Geography. *Journal of Political Economy*, 99(3), 483–499.
- Redding, S. and Rossi-Hansberg, E. (2017). Quantitative Spatial Economics. *Annual Review of Economics*, 9, 21–58.