

Macro-Spatial Economics

Lecture 6: Labour Mobility, Local Shocks, and the Limits of Spatial Arbitrage

Roberto Pancrazi

University of Warwick

Where We Are

The thread so far:

- **Lecture 1** (HM19): spatial misallocation causes large aggregate welfare losses
- **Lectures 2–4** (AA): the full quantitative spatial GE toolkit — gravity, hat algebra, welfare accounting
- **Lecture 5** (DR13, HM24): aggregate losses modest when equilibrium works; workers keep less than expected; 38% of gains travel via migration

Today: the migration margin that propagates 38% of gains — how does it actually work, and what determines its size?


Two questions

MRR (2018): why does the same-sized demand shock move employment by 0.5 in one county and 2.5 in another?

Notowidigdo (2020): why do low-skill workers not leave declining cities the same way high-skill workers do?

The Empirical Puzzle: Heterogeneous Employment Responses

Consider two counties hit by a comparable Bartik demand shock of equal magnitude. One sees employment rise 0.5%; the other 2.5%.



- Mean elasticity ≈ 1.52
- Range: 0.5 to 2.5
- Standard practice: estimate one elasticity, apply everywhere
- **Problem:** that elasticity can be wrong by a factor of 5 for specific counties

The local employment elasticity is not a structural parameter. It is an *endogenous outcome* of the model — driven by how connected a county is to its neighbours through commuting and trade.

The Distributional Puzzle: Low-Skill Workers Stay

When a city experiences an adverse demand shock, its college-educated share tends to *fall* (Glaeser & Gyourko 2005). Low-skill workers leave more slowly.

Standard explanation: higher mobility costs for low-skill workers.

Notowidigdo's alternative: low-skill workers face a *lower incidence* of adverse shocks.

- Housing costs fall when cities shrink — and low-skill workers spend more of income on housing
- Means-tested transfers (food stamps, income maintenance) partially compensate low-skill workers
- So out-migration is less urgent — not because moving is harder, but because staying hurts less

The welfare state is not an afterthought. It is a structural parameter of the spatial equilibrium, and it determines who moves and who stays.

The Plan

MRR (2018)

Commuting, Migration, and Local Employment Elasticities

- AA + commuting as third spatial friction
- Endogenous employment elasticity
- Gravity for goods *and* commuting flows
- Quasi-experiment: million dollar plants

Notowidigdo (2020)

The Incidence of Local Labour Demand Shocks

- Two worker types, durable housing
- Spatial equilibrium with transfers
- Asymmetric adjustment to +/- shocks
- Bartik IV, nonlinear GMM

Two dimensions of the same question

MRR: *who* responds and *how much* — cross-sectional heterogeneity driven by commuting structure.

Notowidigdo: *whether* workers move at all — time-series asymmetry driven by housing durability and transfer incidence.

The Model at a Glance: Standard NEG Building Blocks

Production (tradeable varieties, monopolistic competition):

- Firms use labour only; fixed cost F + variable cost scaled by $1/A_i$
- CES demand \Rightarrow equilibrium markup $\sigma/(\sigma - 1)$; measure of varieties $M_i = L_i/(\sigma F)$
- Trade via iceberg costs $d_{ni} \Rightarrow$ gravity for goods (eq. 6)

Land: inelastically supplied (H_n); owned by immobile absentee landlords who consume only where they live (no mechanical externality on worker location).

Workers: endowed with one unit of labour; choose residence–workplace pair (n, i) to maximise utility; Cobb-Douglas over goods (α) and residential land ($1 - \alpha$); Fréchet preference heterogeneity over (n, i) pairs.

Agglomeration forces

- Home-market effects (CES + IRS)
- More varieties \Rightarrow lower P_n

Dispersion forces

- Inelastic land \Rightarrow rising Q_n
- Commuting costs κ_{ni} , trade costs d_{ni}

Setup: The Key Departure from AA

Economy: locations $n, i \in \mathcal{N}$, land endowment H_n , measure \bar{L} workers.

Each worker chooses a residence-workplace pair (n, i) — not a single location.

This is the key modelling decision. In AA, workers pick a location n (migrate). In MRR, workers pick *where to live* and *where to work* separately. Commuting allows these to differ.

Consequences:

- Employment $L_n \neq$ residents R_n in general
- High-productivity locations can attract employment without attracting residents
- The local labour supply elasticity depends on the *commuting* market as well as the migration market

MRR \supset AA: if commuting costs are prohibitive, workers must live where they work and the model collapses to AA.

Preferences: Fréchet over Commuting Pairs

Utility of worker ω who lives in n and works in i :

$$U_{ni}^{\omega} = \frac{b_{ni}^{\omega}}{\kappa_{ni}} \left(\frac{C_n^{\omega}}{\alpha} \right)^{\alpha} \left(\frac{H_n^{\omega}}{1-\alpha} \right)^{1-\alpha} \quad (1)$$

where $\kappa_{ni} \in [1, \infty)$ is the iceberg commuting cost (utility cost of the commute). C_n^{ω} and H_n^{ω} are consumption and residential land in the *residence* location n (not the workplace i).

Idiosyncratic amenities b_{ni}^{ω} : drawn from Fréchet:

$$G_{ni}(b) = e^{-B_{ni} b^{-\epsilon}}, \quad B_{ni} > 0, \epsilon > 1 \quad (2)$$

B_{ni} = average amenity for commute pair (n, i) ; ϵ = preference dispersion.

Why Fréchet? Same reason as in EK (trade) and AA (migration):

- Max of Fréchet is Fréchet \Rightarrow closed-form choice probabilities
- Generates a gravity equation for commuting flows
- ϵ plays the same role as Fréchet shape in AA: higher ϵ = lower preference dispersion = more elastic labour supply

Commuting Gravity

The probability that worker ω lives in n and works in i (eq. 10):

$$\lambda_{ni} = \frac{B_{ni} (\kappa_{ni} P_n^\alpha Q_n^{1-\alpha})^{-\epsilon} w_i^\epsilon}{\sum_{r \in \mathcal{N}} \sum_{s \in \mathcal{N}} B_{rs} (\kappa_{rs} P_r^\alpha Q_r^{1-\alpha})^{-\epsilon} w_s^\epsilon} \equiv \frac{\Phi_{ni}}{\Phi} \quad (10)$$

The conditional commuting share — probability of working in i given residing in n :

$$\lambda_{ni}^R = \frac{B_{ni} (w_i / \kappa_{ni})^\epsilon}{\sum_{s \in \mathcal{N}} B_{ns} (w_s / \kappa_{ns})^\epsilon} \quad (12)$$

This is a gravity equation for commuting

- Bilateral resistance: w_i, B_{ni}, κ_{ni} in the numerator
- Multilateral resistance: all alternative workplaces s in the denominator
- Elasticity of commuting with respect to commuting costs: $-\epsilon$

Structurally identical to AA's migration gravity — commuting costs κ_{ni} play the role of migration frictions.

Migration and the Two-Stage Structure

Workers choose residence first; then workplace conditional on residence.

Residential choice probability — probability of living in n :

$$\lambda_n^R = \sum_{i \in \mathcal{N}} \lambda_{ni} = \frac{\sum_{i \in \mathcal{N}} \phi_{ni}}{\Phi}, \quad \Phi = \sum_r \sum_s \phi_{rs} \quad (11)$$

Average income of a resident of n (wages weighted by commuting probabilities):

$$\bar{v}_n = \sum_{i \in \mathcal{N}} \lambda_{ni|n}^R w_i \quad (14)$$

\bar{v}_n is high when n has low commuting costs to high-wage workplaces — a *commuting market access* term, exactly parallel to the goods-market access price index in AA.

The two-stage structure means that a residence location's attractiveness depends not just on its own amenities and land costs, but on its commuting access to high-wage employment locations nearby.

Goods Market and Nesting AA

Goods market: standard AA-style CES trade, iceberg costs d_{ni} . Trade share:

$$\pi_{ni} = \frac{L_i (d_{ni} w_i / A_i)^{1-\sigma}}{\sum_{k \in \mathcal{N}} L_k (d_{nk} w_k / A_k)^{1-\sigma}} \quad (6)$$

Nesting: MRR adds commuting to the AA two-friction world (trade + migration).

Model	Trade	Migration	Commuting
AA (2014)	✓	✓	×
MRR (2018)	✓	✓	✓

Commuting costs κ_{ni} and trade costs d_{ni} enter structurally identically — both generate gravity equations calibrated from observed bilateral flows. The MRR gravity equation for commuting is the EK gravity equation with workers instead of goods.

Equilibrium and Hat Algebra

General equilibrium determined by 6 sets of equations: income-expenditure (7), average residential income (14), land market clearing (5), workplace probabilities (11, L_n), residential probabilities (11, R_n), and goods price indices (8).

Free mobility: expected utility equalised across all residence-workplace pairs:

$$\bar{U} = \Gamma \left(\frac{\epsilon - 1}{\epsilon} \right) \left[\sum_r \sum_s B_{rs} (\kappa_{rs} P_r^\alpha Q_r^{1-\alpha})^{-\epsilon} w_s^\epsilon \right]^{1/\epsilon} \quad (15)$$

Counterfactuals without observing fundamentals: rewrite system in terms of proportional changes (hat algebra). Given observed ($L_n, R_n, w_n, \bar{v}_n, \pi_{nn}, \lambda_{ni}$), counterfactual equilibrium solves in terms of $\hat{A}_n, \hat{\kappa}_{ni}, \hat{d}_{ni}$ — exactly as in AA.

The model is methodologically identical to AA: quantify from observed data, solve counterfactuals without measuring unobservable fundamentals.

Measuring Commuting Flows: Data and Gravity Regression

Data: Census 2000 CTPP (Census Transportation Planning Package), county-to-county worker flows, continental US. Unit of observation: county pair (n, i) .

Observable: $\lambda_{ni|n}^R$ — the conditional commuting share (eq. 12): fraction of residents of n who work in i .

From eq. (12), taking logs and parameterising $\kappa_{ni} = \text{dist}_{ni}^{\phi/\epsilon}$ and $B_{ni} = B_n B_i$:

$$\ln \lambda_{ni|n}^R = \delta_n^R + \delta_i^L - \phi \ln \text{dist}_{ni} + \epsilon_{ni}$$

- δ_n^R : residence fixed effect — absorbs P_n , Q_n , residence amenity B_n
- δ_i^L : workplace fixed effect — absorbs w_i^ϵ and workplace amenity B_i
- ϕ : identified *solely* from bilateral distance variation across county pairs

Step 2: recovering ϵ

In a second step, replace the workplace FE with the structural term $\epsilon \ln w_i$. The coefficient on $\ln w_i$ identifies ϵ directly. Result: $\hat{\phi} = 4.43$, $\hat{\epsilon} = 3.30$.

Quantification: Distribution of Employment Elasticities

(MRR Section III)

Procedure: shock each of the 3,111 US counties with a 5% productivity increase (holding all other counties fixed). For each treated county n , compute:

$$\eta_n = \frac{\partial \ln L_n}{\partial \ln A_n}$$

Finding (Figure 2):

- Mean elasticity: $\bar{\eta} \approx 1.52$
- Range: **0.5 to 2.5** — threefold variation
- Resident elasticity (dashed): much less dispersed (0.2–1.2)
- Gap between employment and residents is driven entirely by commuting

Key driver (Table 2)

Regressing $\hat{\eta}_n$ on the residence own-commuting share $\lambda_{nn|n}^R$:

- $\uparrow \lambda_{nn|n}^R$ (closed market) $\Rightarrow \downarrow \eta_n$
- An average elasticity from cross-section data is misleading: applying $\bar{\eta}$ to any individual county can be badly wrong

Robustness: Housing Supply Elasticity (MRR Figure 3)

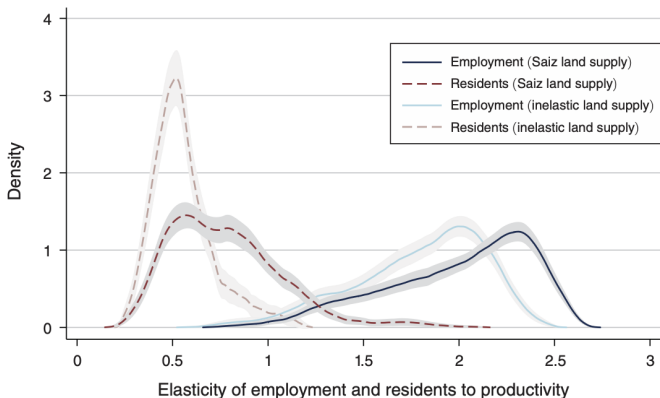
Question: does the heterogeneity in η_n disappear once we allow housing supply to differ across cities?

Exercise: replace inelastic land with Saiz (2010) housing supply elasticities. Saiz sample: 95 MSAs, elasticities from 0.76 (San Jose) to 3.09 (Charlotte).

Robustness: Housing Supply Elasticity (MRR Figure 3)

Question: does the heterogeneity in η_n disappear once we allow housing supply to differ across cities?

Exercise: replace inelastic land with Saiz (2010) housing supply elasticities. Saiz sample: 95 MSAs, elasticities from 0.76 (San Jose) to 3.09 (Charlotte).



Robustness: Housing Supply Elasticity (MRR Figure 3)

Question: does the heterogeneity in η_n disappear once we allow housing supply to differ across cities?

Exercise: replace inelastic land with Saiz (2010) housing supply elasticities. Saiz sample: 95 MSAs, elasticities from 0.76 (San Jose) to 3.09 (Charlotte).

Result (Figure 3):

- Both distributions *shift right*: more elastic housing \Rightarrow larger employment responses
- **Residents** elasticity: heterogeneity *grows* (housing supply matters for who can live where)
- **Employment** elasticity: heterogeneity *persists unchanged* — commuting openness is still the main driver

The key split

Housing supply \rightarrow *residents* (who can live in a city)

Commuting \rightarrow *employment* (who can work in a city)

Connection to HM (2019)

HM focus on housing supply as the binding constraint. MRR: housing matters for *residents*; commuting matters for *employment*.

MDP: Design and Event Study

GHM (2010) design: 82 plant-location competitions from *Site Selection* journal. Each article names the winner county and one or two runner-up counties. Runner-ups survived a long selection process but narrowly lost \Rightarrow valid counterfactual for winners, conditional on controls.

Event-study specification:

$$\ln L_{it} = \kappa l_{j\tau} + \sum_{\tau=-10}^{10} \theta_{\tau} (T_{\tau} \times W_i) + \alpha_i + \eta_j + \mu_t + \varepsilon_{it} \quad (19)$$

i : county; j : case (winner + runner-up group); t : calendar year; τ : years relative to MDP announcement

$l_{j\tau}$: indicator = 1 for case j from treatment year onward; T_{τ} : indicator = 1 for year τ
 W_i : indicator = 1 for winner county; α_i, η_j, μ_t : county, case, and year fixed effects

κ : average employment change in *all* counties after announcement; θ_{τ} : difference-in-differences coefficients

Finding

No pre-trends ($\hat{\theta}_{\tau} \approx 0$ for $\tau < 0$). By 10 years after announcement, employment in winner counties is $\approx 4\%$ **larger** than runner-up counties.

MDP: The Structural Test

MRR's question: not just whether employment rises in winners (it does), but whether the *size of the gain* depends on commuting openness — as the model predicts.

Augmented specification:

$$\ln L_{it} = \kappa I_{jT} + \theta (I_{jT} \times W_i) + \beta (I_{jT} \times \lambda_{ii|i}^R) + \gamma (I_{jT} \times W_i \times \lambda_{ii|i}^R) + \alpha_i + \eta_j + \mu_t + \varepsilon_{it} \quad (20)$$

$\lambda_{ii|i}^R$: residence own-commuting share in 1990 (from CTPP) θ : employment gain for winner with

$\lambda_{ii|i}^R = 0$ (fully *open* local labor market)

β : avg employment change in *both* winner and runner-up counties (placebo check)

γ : differential employment gain for winners as market closes

Model predictions

- $\theta > 0$: open markets see the largest gains
- $\beta \approx 0$: no shock in runner-up counties: commuting openness does not matter
- $\gamma < 0$: gain falls with closeness

Why this matters

If the model's commuting mechanism is correct, $\gamma < 0$ should hold. This is a **structural prediction**, not a reduced-form average — it would not arise in a model without commuting.

MDP: Results (MRR Table 3)

Coefficient	Dep. var.: $\ln L_{it}$			
	(1) Baseline	(2) Res. share	(3) Work. share	(4) Average
$\theta: I_{j\tau} \times W_i$	0.057***	0.250***	0.191***	0.244***
$\beta: I_{j\tau} \times \lambda_{ii i}^{(\cdot)}$	—	0.012	0.243*	0.124
$\gamma: I_{j\tau} \times W_i \times \lambda_{ii i}^{(\cdot)}$	—	-0.242***	-0.177**	-0.241***

- **Column (1):** baseline DiD. A MDP raises county employment by **5.7%** on average — consistent with GHM.
- **Column (2):** $\hat{\gamma} = -0.242 < 0$ and significant. Winner counties with more open commuting markets experience **larger** employment gains. $\hat{\beta} \approx 0$: placebo passes — runner-up counties show no differential effect of commuting openness.
- **Columns (3)–(4):** robust to using workplace or average commuting share instead of residence share.

Conclusion

The quasi-experiment validates the model's structural prediction: commuting openness ($\lambda_{ii|i}^R$) determines the size of the employment response to a local labor demand shock. This is the mechanism behind the heterogeneous η_n in Figure 2.

Counterfactual: Reducing Commuting Costs

Now use the validated mechanism for policy. Exercise: reduce commuting costs by the median observed decline over 1990–2010 (-12%).

Welfare result: $+3.3\%$ — comparable to eliminating substantial trade barriers.

Spatial reallocation pattern:

- **Employment** concentrates further in net-importer counties (e.g. Manhattan): lower commuting costs allow more workers to access high-wage workplaces without moving residence
- **Residents** cluster in net-exporter counties (e.g. New Jersey suburbs): lower costs make it cheaper to live in pleasant locations and commute in

Interpretation

Commuting improvements are a **substitute for housing**

liberalisation: both allow workers to access high-productivity locations. The 3.3% welfare gain comes without anyone having to move home.

Connection to HM (2019)

HM ask: what is the cost of housing constraints blocking migration? MRR answer: commuting can partly offset those constraints — but only within commutable distance.

What MRR Cannot Tell Us

MRR answers: why does the employment elasticity range from 0.5 to 2.5 across counties?

Answer: commuting structure.

But MRR is silent on three things:

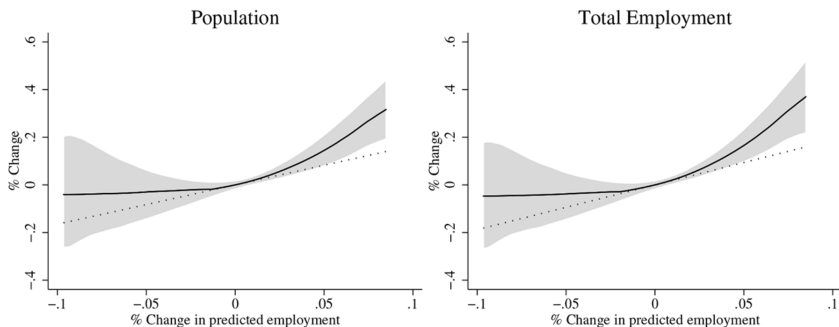
- 1 **Sign asymmetry:** why do cities adjust more to positive than negative shocks?
- 2 **Skill composition:** why do low-skill workers leave declining cities more slowly than high-skill workers?
- 3 **Institutions:** MRR has no transfers, no welfare state, no housing durability

Notowidigdo's contribution

The same spatial equilibrium framework, extended with two ingredients: *durable housing* (supply is more elastic upward than downward) and *means-tested transfers* (low-skill workers receive insurance against adverse shocks). Both change the no-arbitrage condition — and together they explain the patterns MRR cannot.

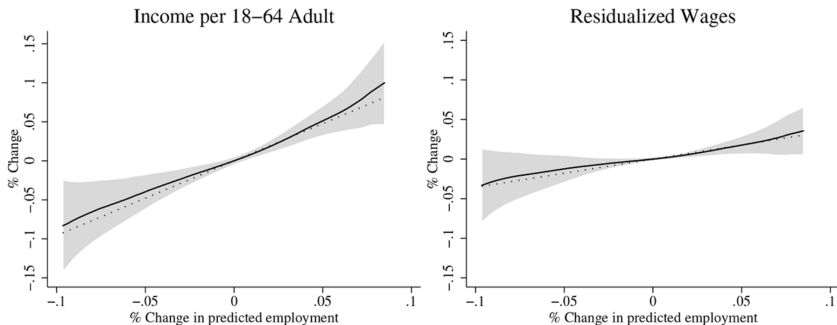
The Asymmetry in the Data

Key empirical finding: positive Bartik shocks increase population *more* than equal-sized negative shocks reduce population.



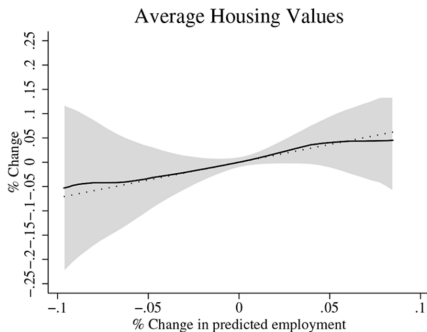
The Asymmetry in the Data

Not for income



The Asymmetry in the Data

Not for housing prices



Reading the Data: A Process of Elimination

Fact: population and employment respond *asymmetrically*. What is the cause?

Moment	Rules out	Why
Wages symmetric	Price/wage rigidity	Rigid wages \Rightarrow wages <i>also</i> asymmetric
House prices symmetric	Mobility frictions	Immobile workers keep housing demand up in declining cities \Rightarrow prices asymmetric

What remains: concave housing supply (durable housing)

Cities build when demand rises; cannot demolish when demand falls.

\Rightarrow Population asymmetric; prices symmetric. \checkmark

Model Setup

Two periods, two worker types $j \in \{H, L\}$ (high-skill, low-skill). Single city receiving a shock against a large outside economy.

Production: CES aggregate,

$$y_{it} = \theta_{it} \left((1 - \lambda)L_{it}^\rho + \lambda(\zeta H_{it})^\rho \right)^{\alpha/\rho}$$

θ_{it} = city-level labour demand shock; ζ = relative high-skill efficiency; ρ = substitution parameter.

Wages equal marginal products. Totally differentiating:

$$\Delta w_{it}^H = \Delta \theta_{it} + ((\rho - 1) + (\alpha - \rho)\pi) \Delta H_{it} + (\alpha - \rho)(1 - \pi) \Delta L_{it} \quad (1)$$

$$\Delta w_{it}^L = \Delta \theta_{it} + ((\rho - 1) + (\alpha - \rho)(1 - \pi)) \Delta L_{it} + (\alpha - \rho)\pi \Delta H_{it} \quad (2)$$

where $\pi = \lambda(\zeta H)^\rho / ((1 - \lambda)L^\rho + \lambda(\zeta H)^\rho)$ is the high-skill share of the labour aggregate.

Transfer Payments in the Model

Means-tested transfers available *only* to low-skill workers:

$$b_{it} = \bar{B} (w_{it}^L)^\Psi$$

\bar{B} = generosity constant; Ψ = elasticity of transfers w.r.t. low-skill wages. Totally differentiating:

$$\Delta b_{it} = \Psi \Delta w_{it}^L, \quad \Psi < 0 \quad (3)$$

$\Psi < 0$: transfers rise when wages fall (wage insurance).

Let $s_b^L = b^L / (w^L + b^L)$ = transfers as share of total income for low-skill; $s_b^H = 0$ by assumption. Let $s_{\mathcal{H}}^j$ = housing expenditure share for skill j (with $s_{\mathcal{H}}^L > s_{\mathcal{H}}^H$ in the data).

Why this matters structurally

When an adverse shock hits, $\Delta w^L < 0 \Rightarrow \Delta b > 0$ (transfers rise). The effective income loss for low-skill workers is attenuated by $(1 - s_b^L)$. This shifts the no-arbitrage condition for low-skill workers: staying hurts less, so less out-migration occurs.

Housing Market: Constant-Elasticity vs. Concave (Durable Housing)

Housing demand: $\mathcal{H}^D(p_{it}^N) = (s_{\mathcal{H}}^H w_{it}^H H_{it} + s_{\mathcal{H}}^L (b_{it}^L + w_{it}^L) L_{it}) / p_{it}^N$. Market clearing:

$$\Delta p_{it}^N + \Delta \mathcal{H}^S(\Delta p_{it}^N) = \nu(\Delta y_{it}^H + \Delta H_{it}) + (1 - \nu)(\Delta y_{it}^L + \Delta L_{it}) \quad (4)$$

Constant elasticity: $\Delta \mathcal{H}^S = \sigma \Delta p_{it}^N$.

- Perfect mobility: symmetric population, wages, prices
- With mobility costs: asymmetric population *and* prices

Concave (durable housing):

$\Delta \mathcal{H}^S(\Delta p^N)$ is convex in Δp^N (units not destroyed — G&G 2005).

- Perfect mobility: prices *symmetric*, population asymmetric
- With mobility costs: both asymmetric

Identification logic

Population asymmetry \Rightarrow concave housing supply and/or mobility costs.

Housing-price asymmetry \Rightarrow specifically mobility costs. The data show population

asymmetry *without* housing-price asymmetry \Rightarrow durable housing is the primary mechanism, not mobility costs.

No-Arbitrage and the Labour Supply Curve

Spatial equilibrium: the marginal worker is indifferent between staying and paying c to out-migrate.

Low-skill no-arbitrage (implicit labour supply curve):

$$(1 - s_b^L) \Delta w_{it}^L + s_b^L \Delta b_{it} - s_{\mathcal{H}}^L \Delta p_{it}^N + c^L(\Delta L_{it}) = 0 \quad (5)$$

High-skill no-arbitrage:

$$\Delta w_{it}^H - s_{\mathcal{H}}^H \Delta p_{it}^N + c^H(\Delta H_{it}) = 0 \quad (6)$$

Three reasons why $\Delta L_{it} \neq \Delta H_{it}$ for equal shocks:

- ① $s_b^L > s_b^H = 0$: transfers cushion low-skill but not high-skill
- ② $s_{\mathcal{H}}^L > s_{\mathcal{H}}^H$: housing price declines compensate low-skill more
- ③ Potentially $c^L \neq c^H$: mobility costs differ by skill

The GMM estimates isolate the relative importance of each mechanism.

Empirical Strategy: Bartik IV, Used Asymmetrically

Standard Bartik (1991) instrument: interact initial industry shares with national industry growth:

$$\Delta \hat{\theta}_{it} = \frac{\hat{E}_{it} - E_{i,t-\tau}}{E_{i,t-\tau}}, \quad \hat{E}_{it} = \left(1 + \sum_k \varphi_{i,k,t-\tau} \frac{u_{-i,k,t} - u_{-i,k,t-\tau}}{u_{-i,k,t-\tau}} \right) E_{i,t-\tau}$$

$\varphi_{i,k,t-\tau}$ = initial employment share of industry k in city i ; $u_{-i,k,t}$ = national employment share of industry k excluding city i .

Notowidigdo's innovation: use Bartik *asymmetrically* — allow separate effects for positive and negative shocks of equal magnitude.

Estimating equation:

$$\Delta x_{it} = \beta \Delta \hat{\theta}_{it} + \delta (\Delta \hat{\theta}_{it})^2 + \alpha_t + \Delta \varepsilon_{it} \quad (7)$$

$\delta \neq 0 \Leftrightarrow$ asymmetric response. Marginal effects at +1SD and -1SD compared.
Data: 215 US MSAs, 1980–2000 (IPUMS + REIS).

Results: The Asymmetry Findings

Key reduced-form results from estimating (7):

Outcome	$\hat{\beta}$	$\hat{\delta}$
Population (18–64)	1.83***	28.0***
Employment	2.08***	32.5***
Income per adult	0.96***	0.38 (insig.)
Residualised wages	0.35***	−0.76 (insig.)
Rental prices	0.84***	−1.01 (insig.)
Housing values	0.71**	−2.76 (insig.)
Food stamps + income maint.	−2.38***	−21.8*

Reading the table: with $\hat{\beta}, \hat{\delta} > 0$ for population, the marginal effect at $+1\sigma$ is $\hat{\beta} + 2\hat{\delta}\sigma \approx 3.78$ while at -1σ it is $\hat{\beta} - 2\hat{\delta}\sigma \approx -0.13$ — positive shocks raise population far more than negative shocks reduce it (convex response).

Population/employment: strongly asymmetric ($\hat{\delta}$ significant). Wages, rents, house values: no asymmetry ($\hat{\delta}$ insignificant). Consistent with concave housing supply + limited mobility costs.

Notowidigdo GMM Estimation

Goal: recover (i) the shape of the housing supply curve, (ii) mobility costs by skill (c^H , c^L), and (iii) housing expenditure shares ($s_{\mathcal{H}}^H$, $s_{\mathcal{H}}^L$).

Method: nonlinear simultaneous-equations GMM. System of equations (1)–(6) solved jointly. Moment conditions exploit the joint distribution of $(\Delta w^j, \Delta L^j, \Delta H^j, \Delta p^H, \Delta b)$ across MSAs and over time.

Identification requirement: unobserved changes in local amenities induced by labor demand shocks are *not differentially valued by skill*. Stronger assumption than reduced-form.

Key GMM findings:

- Housing supply curve is *concave* (confirmed)
- Mobility costs: modest and comparable for high-skill and low-skill workers
- **Primary explanation for low-skill immobility:** transfers, not mobility costs
- Much of the incidence of adverse shocks diffuses to homeowners, landlords, and public assistance programmes

Notowidigdo Counterfactual: Transfers vs. Mobility Subsidies

Counterfactual: replace means-tested transfers with unconditional mobility subsidies for both skill groups.

Effects:

- Low-skill workers become more mobile: skill composition of declining cities improves
- Total population adjustment remains asymmetric — housing supply curve is still concave
- Mobility of high-skill workers changes little (they received few transfers anyway)

Policy lesson

Mobility subsidies and housing supply reform address *different problems*. Mobility subsidies improve skill sorting and address the insurance distortion. Housing supply reform addresses the concavity of the supply curve. Neither alone eliminates the asymmetry in adjustment — both structural features must be addressed together.

Transfers buy insurance for low-skill workers against adverse local shocks; they do not correct the structural asymmetry from durable housing.

Synthesis: The Full Geometry of Spatial Adjustment

MRR (2018)

- Cross-sectional heterogeneity in η_n
- Driven by commuting structure
- Range: 0.5 to 2.5 across counties
- Policy: commuting \leftrightarrow housing as substitutes

Notowidigdo (2020)

- Time-series asymmetry of adjustment
- Driven by housing durability + incidence
- Low-skill stay because shocks hurt them less
- Policy: transfers vs. mobility subsidies

What goes wrong if you use only one paper:

- MRR alone: you correctly predict *how much* a positive shock raises employment in a well-connected county — but you will over-predict how much a negative shock reduces population, and you will attribute low-skill immobility to high mobility costs (wrong)
- Notowidigdo alone: you correctly model the asymmetry and the incidence — but you treat the employment elasticity as homogeneous, applying an average where MRR shows a factor-of-5 spread

Summary: Results and Lessons

Main results:

- 1 **Endogenous employment elasticity:** $\eta_n \in [0.5, 2.5]$, driven by commuting openness ($\lambda_{nn|n}^R$). Housing supply shapes *residents*; commuting shapes *employment*.
- 2 **Asymmetric population response:** population reacts convexly to shocks (positive \gg negative), while prices remain symmetric \Rightarrow concave housing supply, not mobility costs.
- 3 **Incidence over mobility:** low-skill immobility reflects lower shock incidence (via housing shares and transfers), not higher mobility costs.

Methodological takeaways:

- 1 **Heterogeneity matters:** a single Bartik elasticity masks wide variation. Use $\lambda_{nn|n}^R$ or estimate the full spatial model.
- 2 **Institutions are structural:** transfers (s_b^L) enter no-arbitrage directly; ignoring them biases cross-country and policy analysis.
- 3 **Allow for nonlinearity:** test asymmetry (e.g. $\delta(\Delta\hat{\theta})^2$) to uncover structural mechanisms missed by linear models.

Methods Preview: Lecture 7

The paper we study make a large use of spatial data and regressions

Lecture 7 (Spatial Econometrics)

Lecture 7 develops the tools to work with geographically structured data: spatial weight matrices (\mathbf{W}), Moran's I for global clustering, LISA for local hot-spots, and the Spatial Durbin Model (SDM) for regressions with spillovers. All these tools relate to patterns you have already encountered — commuting gravity, market access propagation, etc...