Macro-Spatial Economics

Lecture 5: Welfare, Misallocation, and Incidence

Roberto Pancrazi

University of Warwick

Where We Are and Where We Are Going

Lectures 2–4: AA gave us the toolkit — quantitative spatial equilibrium, gravity, exact hat algebra, welfare accounting.

Today: we use it. The question shifts from *can we measure spatial welfare?* to *what does spatial welfare actually look like empirically?*

Two sub-questions:

1. **How large are the aggregate welfare losses?**
   DR (2013): decompose losses by source — efficiency, amenities, frictions.
2. **Who bears them?**
   HM (2024): when productivity grows locally, what share accrues to workers vs. landowners vs. other cities?

The thread

Lecture 1 (HM19): are there welfare losses? Yes, possibily.
Today (DR13): how large when equilibrium works? Modest — but why?
Today (HM24): who captures the gains? Less than workers expect.

The Three-Wedge Decomposition

HM (2019): welfare losses from misallocation. But *what drives them, and how large are they when the spatial equilibrium is allowed to work?*

DR (2013): city size N_{it} is determined by three wedges:

1. **Efficiency** A_{it}: productive cities attract too few workers
2. **Amenities** γ_{it}: quality-of-life pulls workers away from productive locations
3. **Excessive frictions** g_{it}: government inefficiency inflates costs in large cities

Strategy: set each wedge to its population-weighted average. Solve the new equilibrium. Compare welfare.

Desmet & Rossi-Hansberg (AER 2013)

192 US MSAs (2005–2008) and 212 Chinese cities (2005).

Firms: Technology and the Efficiency Wedge

Production in city i **at time** t (Cobb-Douglas):

$$Y_{it} = A_{it} K_{it}^{\theta} H_{it}^{1-\theta}$$

A_{it} = city TFP; K_{it} = total capital; $H_{it} = N_{it} h_{it}$ = total hours. Lowercase = per capita: $y_{it} = Y_{it}/N_{it}$.

Firm FOCs (competitive factor markets, r_t common across cities):

$$w_{it} = (1-\theta)\frac{Y_{it}}{H_{it}} = (1-\theta)\frac{y_{it}}{h_{it}}, \qquad r_t = \theta\frac{Y_{it}}{K_{it}} = \theta\frac{y_{it}}{k_{it}} \tag{1}$$

Efficiency wedge — directly measurable:

$$A_{it} = \frac{y_{it}}{k_{it}^{\theta} h_{it}^{1-\theta}} \tag{2}$$

No instruments needed

Given data on y_{it}, k_{it}, h_{it} and parameter θ, the efficiency wedge A_{it} is identified by pure accounting.

Agents: Preferences and Budget Constraint

Lifetime utility of a perfectly mobile worker:

$$\sum_{t=0}^{\infty} \beta^t \left[\log c_{it} + \psi \log(1 - h_{it}) + \gamma_i\right]$$

γ_i = time-invariant city amenity; ψ = preference weight on leisure.

Budget constraint: $c_{it} + x_{it} = r_t k_{it} + w_{it} h_{it}(1 - \tau_{it}) - R_{it} - T_{it}$
τ_{it} = labour-tax wedge; R_{it} = land rents; T_{it} = commuting costs.

Steady-state assumption: $x_{it} = \delta k_{it}$. **Simplified BC**:

$$c_{it} = w_{it} h_{it}(1 - \tau_{it}) - R_{it} - T_{it} \tag{3}$$

FOC for labour (combined with firm FOC eq. 1):

$$(1 - \tau_{it}) = \frac{\psi}{(1-\theta)} \frac{c_{it}}{1 - h_{it}} \frac{h_{it}}{y_{it}} \tag{4}$$

Free mobility: utility equalised across all populated cities:

$$\bar{u}_t = \log c_{it} + \psi \log(1 - h_{it}) + \gamma_i \qquad \forall\, i : \; N_{it} > 0 \tag{5}$$

Geography: Monocentric Cities and Average Rents

Setup: circular city, production at centre, workers at distance d on one unit of land each $\Rightarrow$ area $= N_{it}$ $\Rightarrow$ radius $\bar{d}_{it} = (N_{it}/\pi)^{1/2}$.

Commuting cost: linear in distance, $T(d) = \kappa\, d$ ($\kappa =$ cost per mile).

Rent arbitrage: all agents must be indifferent across locations, so $R_{it}(d) + T(d)$ is constant. At the boundary $R_{it}(\bar{d}_{it}) = 0$, hence:

$$R_{it}(d) + \kappa\, d = \kappa\, \bar{d}_{it} \quad \Longrightarrow \quad R_{it}(d) = \kappa\big(\bar{d}_{it} - d\big)$$

Average rent: integrate over city area (element $2\pi d\, \mathrm{d}d$):

$$AR_{it} = \frac{1}{N_{it}} \int_0^{\bar{d}_{it}} \kappa(\bar{d}_{it} - d)\, 2\pi d\, \mathrm{d}d = \frac{\kappa}{3}\left(\frac{N_{it}}{\pi}\right)^{1/2}$$

Since $AR_{it} \propto N_{it}^{1/2}$, taking logs ($\ln AR_{it} = \text{const} + \frac{1}{2}\ln N_{it}$, then invert):

$$\ln N_{it} = \alpha_1 + 2\ln AR_{it} \tag{6}$$

Testable prediction

A city twice as large has average rents $\sqrt{2}$ times higher. Estimated: $\hat{\beta}_6 = 2.14$ ✓ (cannot reject $\beta_6 = 2$).

Government: Infrastructure, Taxes, and Excessive Frictions

City government levies a labour tax τ_{it} to finance commuting infrastructure.

Government expenditure $\propto$ total commuting miles $\times$ government inefficiency g_{it}:

$$G = g_{it}\, h_{it} w_{it}\, \kappa\, \tfrac{2}{3}\pi^{-1/2} N_{it}^{3/2}$$

Government budget constraint (eq. 8):

$$\tau_{it}\, h_{it} N_{it} w_{it} = g_{it}\, h_{it} w_{it}\, \kappa\, \tfrac{2}{3}\pi^{-1/2} N_{it}^{3/2}$$

Solving: $\tau_{it} = g_{it}\, \kappa\, \frac{2}{3}\left(\frac{N_{it}}{\pi}\right)^{1/2}$. Taking logs:

$$\ln \tau_{it} = \alpha_2 + \ln g_{it} + \tfrac{1}{2} \ln N_{it} \tag{10}$$

"Excessive frictions"

Larger cities mechanically require more infrastructure and so have higher τ_{it}. Only the friction *above* what city size predicts identifies g_{it}. A tax alone is not sufficient — congestion from city size is endogenous and must be netted out.

Equilibrium: City Size as an Implicit Function of the Three Wedges

Combining the firm FOC (eq. 1), agent optimality (eq. 4), government BC (eq. 10), and free mobility (eq. 5) yields an *implicit* equation for N_{it} (eq. 11):

$$\bar{u}_t + \text{const} = \underbrace{f\left(\log A_{it}^{\frac{1}{1-\theta}},\quad \underbrace{s(g_{it},\, N_{it})}_{\text{friction/congestion}}\right)}_{\text{net efficiency after congestion}} + \underbrace{\gamma_i}_{\text{amenity}} \tag{11}$$

Three comparative statics (proven from eq. 11):

$$\frac{dN_{it}}{dA_{it}} > 0, \qquad \frac{dN_{it}}{d\gamma_i} > 0, \qquad \frac{dN_{it}}{dg_{it}} < 0 \tag{12}$$

Intuition:

- $\uparrow A_{it}$: higher wages $\Rightarrow$ workers move in until congestion dissipates the gain.
- $\uparrow \gamma_i$: higher utility $\Rightarrow$ city grows until crowding equalises utility.
- $\uparrow g_{it}$: higher tax τ_{it} $\Rightarrow$ city shrinks until net wages compensate.

Empirical Approach: Validating the Model

Before identifying wedges structurally, DR validate the model's comparative statics via a *sequential* log-linear regression system (eqs. 14–17).

1. **Efficiency → city size**: $\ln N_{it} = \alpha_1 + \beta_1 \ln A_{it} + \varepsilon_{1it}$ $\quad \hat{\beta}_1 = 2.10^{***}$ ✓
2. **Efficiency → friction**: $\ln \tau_{it} = \alpha_2 + \beta_2 \ln \hat{N}_{it}(A_{it}) + \varepsilon_{2it}$ $\quad \hat{\beta}_2 = 0.41^{***}$ ✓
3. **All wedges → rents**: $\ln AR_{it} = \alpha_3 + \beta_3 \ln \widetilde{\tau}_{it} + \beta_4 \varepsilon_{1it} + \beta_5 \varepsilon_{2it} + \varepsilon_{3it}$ $\quad \hat{\beta}_3 > 0$, $\hat{\beta}_4 > 0$, $\hat{\beta}_5 < 0$ ✓
4. **Rents → city size**: $\ln N_{it} = \alpha_4 + \beta_6 \ln AR_{it} + \varepsilon_{4it}$ $\quad \hat{\beta}_6 = 2.14$ (pred. $= 2$) ✓

All 7 coefficients match theory in sign and magnitude

The general equilibrium structure of the model is empirically validated *before* it is used for structural identification.

Identification: Recovering the Three Wedges

Given $\{y_{it}, k_{it}, h_{it}, w_{it}, c_{it}, AR_{it}, N_{it}\}$ and calibrated (θ, ψ, κ):

Step 1 — Efficiency (pure accounting, eq. 2):

$$A_{it} = \frac{y_{it}}{k_{it}^{\theta}\, h_{it}^{1-\theta}}$$

Step 2 — Excessive frictions (residual regression, eq. 20):

$$\ln \tau_{it} - \tfrac{1}{2} \ln N_{it} = \alpha_5 + \varepsilon_{5it}, \qquad \varepsilon_{5it} \equiv \ln g_{it}$$

Mean of $\ln g_{it}$ normalised to zero — identifies *excess* friction only.

Step 3 — Amenity (residual from free mobility, eq. 19):

$$\gamma_i \text{ such that model matches } N_{it} \text{ exactly, given } \bar{u}_t = 10$$

Equivalent to: $\gamma_i = \bar{u}_t - \log c_{it} - \psi \log(1 - h_{it})$.

Exactly identified

Three unknowns $(A_{it}, g_{it}, \gamma_i)$, three equilibrium conditions, each uses *distinct* observables. No joint estimation.

Estimation: Data and Calibration

Sample: US MSAs, 2005–2008 (pooled; time dummies included).

Data sources:

- y_{it}: MSA GDP — Bureau of Economic Analysis (BEA)
- k_{it}: capital stock — BEA fixed-assets tables
- h_{it}: average hours — Current Population Survey (CPS)
- w_{it}, c_{it}: wages and consumption — Census / ACS
- AR_{it}: median rents — National Association of Realtors + ACS

Calibrated parameters (McGrattan & Prescott 2010):

- $\psi = 1.4841$: leisure preference (matches aggregate US hours)
- $\theta = 0.3358$: capital share (standard)
- $r = \delta = 0.02$: common interest and depreciation rates
- $\kappa = 0.002$: commuting cost per mile — backed out from eq. (9) and the intercept of eq. (20): $\alpha_5 = \ln(2/3) + \ln \kappa - \frac{1}{2} \ln \pi$

Validation: External Checks on the Recovered Wedges

DR compare each recovered wedge to *external proxies not used in estimation*.

Efficiency $\hat{A}_{it}$:

- Correlation with average wages: $r = 0.79$ ✓
- Correlation with labour productivity: $r = 0.90$ ✓

Amenity $\hat{\gamma}_i$ — External quality-of-life index (23 proxies: climate, crime, culture, air quality, ...):

- Sign agreement: **22 of 23** (18 significant at 10%) ✓
- High $\hat{\gamma}$: San Diego, San Francisco, Honolulu (coastal, mild climate)
- Low $\hat{\gamma}$: Detroit, Pittsburgh, Cleveland (workers need a wage premium)

Friction $\hat{\tau}_{it}$ (taxes, commuting, government expenditure, ...):

- Sign agreement: 10 of 11 (8 significant at 5%) ✓

The test of urban accounting

All three wedges pass external validation. Passing 22/23 and 10/11 sign tests on *unused* indices is strong evidence each wedge measures the intended object.

US Counterfactuals

Set each wedge to its population-weighted average; solve new equilibrium.

Equalise	Welfare gain	Population reallocation
Efficiency A_{it}	$+1.2\%$	$\approx 37\%$
Amenities γ_{it}	$+0.2\%$	$\approx 20\%$
Frictions g_{it}	$+0.9\%$	$\approx 44\%$
All three	$+1.54\%$	—

The paradox: 20–44% population reallocation yields only 0.2–1.2% welfare gain.
Analogy: Lucas (1987) — eliminating business cycles yields trivial welfare gains because aggregate risk is small when agents are identical and markets are complete. Same logic applies across space.

The spatial equilibrium dissipates rents as workers move in. HM (2019)'s large numbers rely on frictions *blocking* this mechanism.

China: The Payoff of Part 1

Same methodology, 212 Chinese cities (2005).

US welfare gains:

- Efficiency: $+1.2\%$
- Amenities: $+0.2\%$
- Frictions: $+0.9\%$

China welfare gains:

- Efficiency: $+\mathbf{47}\%$
- Amenities: $+\mathbf{13}\%$
- Frictions: -1.5%

Why an order of magnitude larger?

- *Hukou* restrictions keep efficient cities artificially small
- Spatial equilibrium is *blocked* — efficiency advantages cannot be exploited

The spatial equilibrium is a dissipation mechanism. When it functions, aggregate losses are modest. When it is blocked, the stakes are enormous. This motivates Part 2: what happens when the equilibrium *does* function — but workers still capture less than expected?

Bridge: From Aggregate to Incidence

DR's central finding

Aggregate welfare losses from spatial misallocation are modest when the equilibrium functions. The spatial equilibrium is a dissipation mechanism.

The twist: that same dissipation mechanism also dissipates *individual* gains.

When a productive city expands, workers flow in, rents rise, real wages fall back toward the outside option — the spatial equilibrium equalises utility, not wages. DR's representative-agent framework averages over this: the aggregate welfare calculation hides the distributional question entirely.

DR: *how large is the pie lost to misallocation?*
HM (2024): *when the pie expands, who gets the extra slice — and does anyone?*

HM (2024): Overview and Course Connection

HM (2019) implicit assumption: deregulating housing raises output, and workers capture the gains through higher real wages.

HM (2024) challenge: productivity growth raises rents nearly as much as wages. Workers who rent capture only a fraction of the nominal gain.

Ricardo's concern: gains from productivity $\Rightarrow$ inelastic factor (land) captures the surplus. Workers migrate until real wages equalise — migration dissipates the gain for workers, landowners keep it.

Hornbeck & Moretti (REStat 2024)

193 US MSAs, 1980–2010. Manufacturing TFPR shocks to estimate the magnitude and distribution of gains — direct effects in the shocked city, indirect effects elsewhere via migration.

The Model: Setup (Appendix B)

Two cities a and b; fixed total workforce. Production in city c:

$$\ln Y_c = A_c + (1-h)\, n_c$$

A_c = log TFPR; n_c = log employment share; h = housing/land share (so $1-h$ = labour share).

Indirect utility of worker i in city c:

$$v_{ic} = w_c - \beta r_c + x_c + e_{ic}$$

w_c = log wage; r_c = log housing cost; β = housing budget share; e_{ic} = idiosyncratic preference; $e_{ib} - e_{ia} \sim U[-s, s]$.

Housing supply (reduced form):

$$r_c = k_c\, n_c$$

- s: governs mobility — large s = strong location attachment = inelastic local labour supply
- k_c: inverse housing supply elasticity — large k_c = constrained city (e.g. San Francisco)

Direct Effects: Incidence Formulas (Eqs. 5–8)

City b receives shock Δ; cities initially identical. Equilibrium changes:

$$n_{b2} - n_{b1} = \frac{1}{\beta(k_a + k_b) + 2h + s}\Delta \tag{5}$$

$$w_{b2} - w_{b1} = \frac{\beta(k_a + k_b) + h + s}{\beta(k_a + k_b) + 2h + s}\Delta \tag{6}$$

$$r_{b2} - r_{b1} = \frac{k_b}{\beta(k_a + k_b) + 2h + s}\Delta \tag{7}$$

Purchasing power (real wage gain) in city b:

$$(w_{b2} - w_{b1}) - \beta(r_{b2} - r_{b1}) = \frac{\beta k_a + h + s}{\beta(k_a + k_b) + 2h + s}\Delta \tag{8}$$

- Large s (immobile workers) $\Rightarrow$ workers capture more
- Large k_b (inelastic housing) $\Rightarrow$ landowners capture more; if $k_b \to \infty$, entire Δ capitalised into rents — workers gain nothing

Indirect Effects and Connection to AA (Eq. 9)

City a (no direct shock): out-migration $\Rightarrow$ real wages rise:

$$(w_{a2} - w_{a1}) - \beta(r_{a2} - r_{a1}) = \frac{\beta k_a + h}{\beta(k_a + k_b) + 2h + s}\,\Delta \quad (9)$$

Indirect gain (9) $<$ direct gain (8) unless $s = 0$ (perfect mobility).

HM Appendix B

- 2 cities, no goods trade
- Uniform e_{ic}: closed forms
- Housing: $r_c = k_c n_c$

Allen & Arkolakis

- N locations, gravity trade
- Fréchet: gravity in migration
- Congestion: $u_i = \bar{u}_i L_i^{\beta}$

HM's model is the minimal spatial equilibrium with closed-form incidence. AA generalises to N locations with trade and exact welfare accounting — at the cost of closed forms.

From Theory to Data

The model tells us *what to look for*: a TFPR shock Δ should raise employment, wages, and rents in the shocked city, with the split between workers and landowners governed by housing supply (k_c), mobility (s), and labour demand curvature (h).

Three empirical challenges:

1. **Measuring the shock**: what is city-level TFPR, and how do we estimate it?
2. **Causal identification**: TFPR is endogenous — productive cities attract workers, which raises output and measured TFPR. We need an instrument.
3. **Magnitude and incidence**: how large are the effects on wages vs. rents, and who captures the surplus — workers or landowners?

The next slides address each challenge in order: measurement (TFPR regression), identification (shift-share IV), and results (direct and indirect effects).

TFPR: The Plant-Level Regression (Appendix A)

Separate OLS for each decade (CMF), plant f, city c, decade d:

$$\ln Y_{fcd} = \hat{\beta}_1 \ln K_{fcd} + \hat{\beta}_2 \ln L_{fcd} + \hat{\beta}_3 \ln M_{fcd} + \phi_{cd} + \varepsilon_{fcd}$$

- $\hat{\beta}_1, \hat{\beta}_2, \hat{\beta}_3 =$ capital, labour, materials elasticities (weighted by plant output)
- $\phi_{cd} =$ city fixed effect *in decade d*; $\hat{\phi}_{cd} \equiv \ln \hat{A}_{cd}$ (log TFPR)
- TFPR shock: $\Delta \ln \hat{A}_c = \hat{\phi}_{c,1990} - \hat{\phi}_{c,1980}$

City-level TFPR $= \hat{\phi}_{cd}$ directly from the regression. No industry-level aggregation step.

Data: Census of Manufactures (CMF) — ≈50,000 plants per year; capital, labour, materials inputs, and revenue. Manufacturing chosen because no other sector provides comparable input detail.

Revenue Productivity, Not Physical TFP

Output Y_{fjct} = revenue deflated by an *industry* price index, not plant-specific prices.

TFPR $\neq$ TFPQ (physical TFP):

- $\hat{\phi}_c$ = physical productivity + local output price wedge
- A city with expensive housing commands higher nominal wages and output prices $\Rightarrow$ higher TFPR without more efficient production

Why this is acceptable here: HM (2024) study incidence of productivity growth. The relevant object is the revenue gain that firms and workers respond to — TFPR captures this.

Connects to DR's efficiency wedge A_{it}: both measure effective returns to labour, not physical units produced. Neither is pure technical efficiency.

Measurement Error and the IV Motivation

$\hat{\phi}_c$ estimated from finite plant samples $\Rightarrow$ classical measurement error:

$$\Delta \ln \hat{A}_c = \Delta \ln A_c + u_c, \qquad u_c \perp \Delta \ln A_c$$

OLS problem: attenuation bias — estimated effects on employment and earnings are understated. Also: reverse causality (growing cities attract productive firms).

Solution — shift-share IV:

$$\Delta \ln \hat{A}_c^{IV} = \sum_j s_{jc} \, \Delta \ln \hat{A}_j^{-c}$$

s_{jc} = city c's *output* share in industry j (not employment share). $\Delta \ln \hat{A}_j^{-c}$ = national TFPR growth in industry j computed *excluding city* c's plants — so the instrument is not contaminated by city c's own productivity trend.

- Identification: city's industry mix was set before the shock (1980); national trends are exogenous to any one city
- The leave-out mean is the key Goldsmith-Pinkham requirement: without it, the "national" trend reflects city c itself
- Three additional IVs (stock returns, exports, patents): overidentification tests pass

Data and Outcomes

Sample: 193 US MSAs, 1980–2010.

TFPR shock: 1980–1990 change in $\ln \hat{A}_c$ from CMF. Shock held fixed; outcomes measured at three horizons.

Outcomes (US Census / ACS):

- **Employment**: total and by education group
- **Earnings**: nominal wage per worker
- **Rents and home values**: from Census housing module
- **Purchasing power**: earnings minus housing cost share — separately for renters and homeowners

Controls: Census region fixed effects α_r; baseline city characteristics. Robust standard errors (heteroskedasticity-consistent).

Empirical Specification

Regress changes in outcome Y_c on the 1980–1990 TFPR shock at three horizons:

$$\Delta_{90} \ln Y_c = \pi^M \Delta_{90} \ln \hat{A}_c + \alpha_r + \varepsilon_c \quad (1)$$

$$\Delta_{00} \ln Y_c = \pi^L \Delta_{90} \ln \hat{A}_c + \alpha_r + \varepsilon_c \quad (2)$$

$$\Delta_{10} \ln Y_c = \pi^{XL} \Delta_{90} \ln \hat{A}_c + \alpha_r + \varepsilon_c \quad (3)$$

α_r = Census region FE; TFPR shock always from 1980–1990.

Four IVs — all shift-share: national shock $\times$ initial industry share:

- **Baseline**: leave-out national TFPR growth by industry
- **Stock market**: industry equity returns 1980–1990
- **Export**: export growth from other high-income countries
- **Patent**: patenting activity by technology class

Overidentification tests: all four IVs give statistically identical estimates.

Direct Effects: Employment, Earnings, and Rents

1% increase in city TFPR (1980–1990) — Table 2, Baseline IV:

	Medium run	Long run	Longer run
Employment	+2.38%	+4.16%	+4.03%
Earnings	+0.91%	+1.45%	+1.46%
Rent	+0.98%	+1.47%	+1.09%
Home value	+1.74%	+2.46%	+3.05%

The key compression

Earnings +1.45%, rent +1.47% — nearly equal in the long run. Most of the wage gain is wiped out by higher housing costs for renters.

Why Such Large Employment Effects? The Local Multiplier

Employment rises +4.16% per 1% TFPR. Why?

Manufacturing TFPR growth raises local wages and employment, which increases demand for *nontraded* goods and services (restaurants, retail, construction, healthcare).

Implied multiplier (Appendix Table 6): additional nonmanufacturing jobs per manufacturing job created by TFPR growth:

Horizon	Nonmanufacturing jobs per mfg job
Medium run (1990)	1.62
Long run (2000)	2.21

Consistent with Moretti (2010): each skilled job in tradables generates ≈2 additional local service jobs.

The large employment response is mostly *migration* (in-movers filling new service jobs), not incumbent workers switching sectors. This is why rents rise so much — population grows, not just wages.

Purchasing Power: Renters vs. Homeowners

Long-run purchasing power per 1% TFPR — Table 2, Panel E:
Renters: $+0.62\%$

- Nominal: $+1.45\%$ earnings
- Minus $0.56 \times 1.47\%$ rent (where $0.56 = \underbrace{0.33}_{\text{housing share}} + \underbrace{0.35 \times 0.67}_{\text{nontradable markup}}$)
- $\approx 43\%$ of the nominal wage gain; renters lose roughly two-thirds to higher local costs

Homeowners (insulated from rent): $+1.11\%$
Homeowners (benefits from equity): $+1.60\%$

Cities with inelastic housing supply: rent response $+2.3\%$ vs. $+1.2\%$ in elastic cities.

Ricardo was right: landowners capture a substantial share. Housing tenure determines who benefits — not just wage levels.

Education, Mobility, and Local Inequality

Mechanism: more educated workers are more mobile.

- College: more elastic local supply $\Rightarrow$ smaller wage gains, larger employment response
- HS graduates: inelastic supply $\Rightarrow$ larger wage gains, smaller migration

Long-run estimates per 1% TFPR:

- Employment: college $+5.82\%$ vs. HS $+3.23\%$
- 90–10 earnings gap: -0.998% — productivity growth *compresses* local inequality

Unlike skill-biased technological change, local TFPR growth is locally *equalising* — less mobile workers benefit more at the local level. The intuition: college workers flood in (elastic supply), diluting their own wage gains; HS workers stay put (inelastic supply), capturing a larger share.

Indirect Effects

TFPR shock in city b pulls workers from city a: wages $\uparrow$, rents $\downarrow$ in a.

Three-step estimation:

1. Employment gain in shocked city c: $\Delta L_c = \hat{\pi}^L \times \Delta \ln \hat{A}_c \times L_{c,1980}$
2. Assign workers to origins via observed migration shares s_{oc}: $\Delta L_{o\leftarrow c} = -s_{oc}\,\Delta L_c$
3. Pressure on wages and rents in o: city-level housing supply elasticities from Saiz (2010) — estimated from geographic constraints (water, steep terrain) and regulatory barriers; -0.15 labour demand elasticity for wages

Main finding: **38%** of the combined purchasing power gain accrues *outside* the shocked city.

- College workers: more mobile $\Rightarrow$ disproportionately benefit indirectly
- Renters in origin cities: rents fall — purchasing power rises
- Indirect effects *reverse* the local inequality compression at the aggregate level

The Migration Margin: Looking Ahead to MRR (2018)

The key HM (2024) finding: 38% of the combined purchasing power gain from a local TFPR shock accrues *outside* the shocked city — via out-migration raising wages and lowering rents elsewhere.

Implication: any city-level evaluation of a local shock misses more than a third of the aggregate welfare effect. Gains propagate widely — precisely because workers move.

This raises a deeper question: if workers who move capture less than those who stay (wages rise for stayers, rents fall in origin cities), why do workers sort into productive cities at all — and *who* sorts?

Lectures 6 onward

Monte, Redding & Rossi-Hansberg (2018): how do commuting and migration margins interact, and what are the local employment elasticities? The 38% indirect-effects finding is precisely what motivates that question.

The Two Papers in Dialogue

DR (2013)

- US misallocation cost: modest ($< 2\%$)
- Spatial equilibrium dissipates rents
- China: equilibrium blocked $\Rightarrow$ enormous stakes

HM (2024)

- Productivity gains partly captured by landowners
- Workers keep $\approx$43% of nominal gains
- 38% of gains accrue in other cities

The unifying insight

The spatial equilibrium is a double-edged sword: it limits aggregate welfare losses from misallocation (DR), but also limits workers' gains from productivity growth (HM). Both flow from the same force — free mobility dissipating location rents.

Back to Hsieh & Moretti (2019)

1. **HM (2019)**: housing constraints cause large GDP losses; workers cannot move to productive cities.
2. **DR (2013)**: if workers could move freely, welfare gains are modest (1–2%); equilibrium dissipates the gains.
3. **HM (2024)**: even when workers do move, landowners capture much of the surplus; renters keep ≈one-third of nominal gains.

Whether housing is deregulated or not, workers face a bind. Under current restrictions they cannot relocate (HM 2019); if restrictions lift, rents rise and dissipate the gain (HM 2024). Welfare consequences depend critically on land ownership, not just spatial frictions.

Policy Implications

1. **Housing deregulation**:
   - Raises aggregate output (HM 2019)
   - Net welfare effects modest if equilibrium functions (DR 2013)
   - Renters gain less than aggregate figures suggest (HM 2024)
2. **Place-based investment**:
   - Small welfare returns when equilibrium functions (US)
   - Large returns when mobility is restricted (China)
3. **Land value taxation**:
   - HM (2024) shows a large share of TFPR gains capitalises into land values
   - A tax on land value — not structures — captures this surplus without distorting development
   - Classical Henry George argument with direct empirical backing

Summary

- **DR (2013)**: spatial misallocation costs are modest in equilibrium ($< 2\%$ US); gains from equalising A_{it} swamp gains from equalising γ_{it} or g_{it}
- **HM (2024)**: renters keep $\approx 43\%$ of nominal wage gains; 38% of total gains spill outside the shocked city via migration
- **Open**: what generates spatial TFP gaps? how do gains propagate dynamically? what role for trade linkages?

HM (2019): *are there losses?* → DR (2013): *modest, driven by efficiency* → HM (2024): *who captures the gains?* → **Next**: if gains travel via migration, how does the migration margin actually work — and who moves?