Macro-Spatial Economics

Lecture 9: Dynamic Spatial Models

Roberto Pancrazi

University of Warwick

The Course So Far

What we have established:

- **Lecture 5** (DR13, HM24): who bears the costs of local productivity growth; incidence falls on workers, homeowners, and transfers
- **Lecture 6** (MRR, Notowidigdo): commuting structure determines cross-sectional heterogeneity; housing durability and transfers explain adjustment asymmetry
- **Lecture 8** (M&S 2014, BHO 2019): housing is the primary transmission channel between local shocks and aggregate outcomes; regional variation reveals aggregate dynamics

The gap

All of these frameworks are static or two-period. They compare steady states. They cannot account for transition dynamics, endogenise tenure choice, or explain why welfare-improving housing policies face systematic political opposition.

Four Things Static Models Cannot Do

The state-of-the-art QSM has rich geography and commuting but is static. Abstracting from dynamics obscures four things:

1. **Individual welfare requires transition dynamics**: distributional effects of shocks depend on who is hit, when, and in what tenure state — none of which can be captured in a steady-state comparison
2. **Short-run and long-run effects differ**: the first-best static policy may be infeasible if voters respond to early-period losses
3. **History dependence is ignored**: durable structures and migration frictions mean that the current allocation reflects past shocks; reallocation may be suboptimal even when it appears efficient
4. **Real estate ownership is a forward-looking decision**: welfare from a house price change depends on when the household bought, with what leverage, and whether it can adjust

Lecture 9

Two papers address these failures: **DUE** (Greaney, Parkhomenko & Van Nieuwerburgh 2025) at city scale; **CDP** (Caliendo, Dvorkin & Parro 2019) at national scale.

DUE: Bridging Two Traditions

Quantitative Spatial Models
(ARSW 2015 and extensions)

- Many locations, commuting
- Rich geography, heterogeneous agents

× Static
× No tenure choice
× No forward-looking saving

Macro-Housing Models
(Favilukis et al., Berger et al., KMV)

- Forward-looking households
- Tenure choice, collateral, illiquidity
- Life-cycle dynamics and income risk

× Single location
× No commuting
× No spatial frictions

DUE's contribution

Greaney, Parkhomenko & Van Nieuwerburgh (2025) combine both traditions in a single framework: many locations, commuting frictions, forward-looking agents, tenure choice with realistic constraints, and endogenous housing supply. The resulting model generates qualitatively different welfare conclusions from traditional static QSMs.

What Households Decide

A household in DUE makes four types of decisions:

1. **Where to live**: choose a residential location i from I locations in the city (plus an outside option)
2. **Where to work**: conditional on living in i, choose a workplace j, incurring commuting cost d_{ij}
3. **Housing tenure**: rent or own? If owning, what size house h? Housing services: $\mathbf{h}(h, h^r) = \chi h + h^r$, where $\chi > 1$ captures the non-pecuniary benefit of ownership
4. **Consumption and saving**: choose consumption c and accumulate/decumulate liquid financial wealth b

Why the interaction matters

A household that owns a large house in an expensive location has most of its wealth locked in an illiquid asset. If income falls, it cannot quickly sell and downsize — it must cut consumption instead. A renter can simply move to a cheaper unit. This asymmetry is at the heart of DUE's welfare results.

What Shapes These Choices

Financial frictions:

- Collateral constraint: $b \geq -\phi \, p_{it} \, h$ (LTV ceiling at origination)
- Transaction cost ψ of house value lost every time you sell
- If house prices fall and you are near the constraint, you cannot take on additional debt — but you are not forced to sell either. You are stuck

Moving costs: relocating costs $\mu_{ii'}(a) = \mu_0 + \mu_a \cdot a$ in utility — moving gets more expensive with age, generating realistic declining migration rates over the life cycle

Income risk: log productivity $z(\zeta, a) = \zeta + \bar{z}(a)$, where $\bar{z}(a)$ is a hump-shaped lifecycle profile and ζ follows an AR(1). Good or bad draws affect saving, which affects the ability to buy

Preference shocks: at each decision point, households draw idiosyncratic Gumbel preference shocks ϵ^R (for residential location) and ϵ^W (for workplace). These generate the logit choice probabilities familiar from static QSMs and ensure smooth substitution across locations

Flow utility

$u(c, \mathbf{h}) = (c^{1-\eta} \, \mathbf{h}^{\eta})^{1-\gamma}/(1-\gamma)$; bequest motive $v(b) = \vartheta_0 \, b^{1-\gamma}/(1-\gamma)$

The Timing Problem

In principle, all four decisions happen continuously: at every instant, a household could move, switch jobs, buy or sell a house, and adjust consumption.

If we modelled it that way:

- State vector: (b, h, i, j, ζ, a) — **six** variables
- State space includes I^2 residence-workplace pairs (3,025 for SF; 33,489 for New York)
- Must solve an HJB over this entire space *and* track a six-dimensional cross-sectional distribution through time
- For transition dynamics: **computationally infeasible**

But not all decisions are equally frequent:

- Consumption and saving adjust continuously
- Moving house, changing jobs, buying or selling property are *lumpy, infrequent* decisions

DUE exploits this natural distinction. The model separates high-frequency choices (consumption-saving) from low-frequency choices (location, workplace, housing tenure) using a timing device called **shock ages**.

What Is a Shock Age?

A household's life is divided into annual intervals. The boundaries are called **shock ages**, $\mathcal{A}^s = \{0, 1, \ldots, A-1\}$. Discrete choices occur *only* at shock ages:

At each shock age:

1. Draw new productivity ζ' (AR(1) innovation) and preference shocks ϵ^R, ϵ^W
2. Choose residential location i' (paying moving cost if relocating)
3. Choose workplace j (incurring commuting cost $d_{i'j}$)
4. **Receive the PDV of all wages from j, paid as a lump sum**: $y_t(\zeta, a; j) = \int_t^{t+1} e^{-q(s-t)}(1-\tau_z)\, e^{z(\zeta,a)}\, w_{js}\, \mathrm{d}s$
5. Choose owner-occupied house size h' (subject to collateral constraint and transaction cost)

Between shock ages: location i, house h, and productivity ζ are all **fixed**. The household's only decision is how much to consume and save.

The household enters each interval with state $\Omega = (b,\ h,\ i,\ \zeta,\ a)$ — five variables, **not six**. Workplace j is absent. Why?

Between Shock Ages: The Consumption-Saving Problem

With (i, h, ζ) locked and j already paid out, the household solves a standard HJB:

$$\rho V_t(\Omega) = \max_{c,\, h^r} \; u(c, \mathbf{h}(h, h^r)) + \partial_b V_t(\Omega)\, \dot{b}_t + \partial_a V_t(\Omega) + \partial_t V_t(\Omega)$$

Budget constraint: $\dot{b} = q\, b - c - r_{it}\, h^r - (\delta + \tau_h)\, p_{it}\, h$

The only state variable that *moves* is liquid wealth b. This is a one-dimensional PDE in b (conditional on the other states) — extremely fast to solve using the finite-difference methods of Achdou et al. (2022).

The collateral trap: if b hits the constraint ($b = -\phi\, p_{it}\, h$), further borrowing is prohibited ($\dot{b} \geq 0$). The household cannot sell, refinance, or move until the next shock age. This is what makes housing *illiquid* in a precise sense.

The link between intervals

The value function at the *end* of this interval (just before the next shock age) serves as the boundary condition for the discrete location-workplace-housing choices that follow.

Why Shock Ages? Three Computational Advantages

Advantage 1: State space reduction ($I^2 \to I$)

Workplace j is chosen at shock ages and the entire income stream is capitalised as a lump sum. Between shock ages, j has no relevance: it does not appear in the budget constraint and does not affect any decision. The HJB is solved over I locations, not I^2 pairs. For SF: factor-55 reduction. For New York: factor-183.

Advantage 2: Separation of continuous and discrete problems

Between shock ages: a continuous-time HJB with one moving state (b). At shock ages: discrete logit choices with closed-form probabilities. Neither is hard alone. The original difficulty was solving them simultaneously. Shock ages decompose this into two problems solved *sequentially*.

Advantage 3: Continuous-time FOCs

In continuous time, FOCs hold with equality and are sufficient even with kinks from the collateral constraint. In discrete time, FOCs would be inequalities and the endogenous-gridpoint method would fail. Since wealth does not jump between shock ages, the HJB and KF equations yield sparse (tridiagonal) matrices — extremely fast to solve.

The Cross-Sectional Density and Equilibrium Prices

Equilibrium requires market clearing: wages, rents, and prices are **integrals over the density** $g_t(\Omega)$.

Between shock ages

g_t evolves continuously via Kolmogorov Forward:

$$\partial_t g_t = -\partial_b[\dot{b}_t\, g_t] - \partial_a g_t$$

Wealth drifts, people age — no jumps.

At shock ages

g_t undergoes discrete redistributions:

ζ: productivity shocks reshuffle

i: residential choice moves mass

b: income receipt shifts wealth

h: housing choice redistributes

Key insight: even though individuals jump at shock ages, the *aggregate* density changes continuously. The age distribution is uniform — there is a constant flow of households reaching each shock age at every instant. Discrete events happen to different people at different times and average out.

Consequence for prices

Prices — integrals over g_t — move continuously. No discrete jumps in wages, rents, or house prices. The equilibrium can be solved with standard continuous-time methods. A fully discrete-time model would require a large fixed-point problem at every time step.

Production, Housing Supply, and REITs

Traded-good firms (Ahlfeldt et al. 2015 technology):

$$Y_{jt} = Z_j \, L_{jt}^{\alpha} \, HC_{jt}^{1-\alpha}$$

Z_j = location TFP; L_{jt} = efficiency labour; HC_{jt} = commercial floorspace. Competitive firms pay each factor its marginal product: $w_{jt} = \alpha Z_j L_{jt}^{\alpha-1} HC_{jt}^{1-\alpha}$.

Developers: build residential and commercial floorspace. Construction is *irreversible* ($Y_{mit}^{h} \geq 0$):

$$\dot{H}_{mit} = Y_{mit}^{h} - \delta \, H_{mit}$$

Floorspace price equals construction cost when there is positive construction; otherwise price falls until demand absorbs the existing stock. Irreversibility links directly to Notowidigdo (Lecture 6): cities shrink slowly because structures last.

REITs: own all rental and commercial floorspace. Perfectly diversified across locations. Rents satisfy a no-arbitrage condition: $r_{mit} = (\delta + q + \tau_h) \, p_{mit} - \dot{p}_{mit}$. Households hold fractions of their liquid wealth b in REITs, so aggregate shocks to real estate values affect household wealth through this portfolio channel.

Equilibrium

An **equilibrium** is: value functions $V_t(\Omega)$, a density $g_t(\Omega)$, wages w_{jt}, rents r_{mit}, and prices p_{mit}, such that:

1. **Households optimise**: the HJB (between shock ages) and the discrete-choice logit equations (at shock ages) are satisfied
2. **Density is consistent**: g_t evolves via the KF equation between shock ages and through the discrete redistribution cascade at shock ages — consistent with household decisions
3. **Firms optimise**: wages equal the marginal product of efficiency labour; commercial rents equal the marginal product of floorspace
4. **Labour markets clear**: labour supply (integral of $e^z \pi_t^W g_t$ over Ω) equals labour demand in every location j
5. **Housing markets clear**: residential floorspace demand (integral of housing choices over g_t) equals supply in every location i
6. **No-arbitrage**: rents satisfy $r_{mit} = (\delta + q + \tau_h)\, p_{mit} - \dot{p}_{mit}$ everywhere

A **stationary equilibrium** is one in which all objects are time-invariant.

Solving the Model

The shock-age structure makes the solution algorithm modular:

Step 1 — Guess prices: conjecture time paths for wages $\{w_{jt}\}$, rents $\{r_{mit}\}$, and house prices $\{p_{mit}\}$ in all locations

Step 2 — Solve backwards (HJB): given prices, solve the household problem backwards in age. Between shock ages: finite-difference HJB (one-dimensional in b, fast). At shock ages: evaluate the logit choice probabilities in closed form

Step 3 — Simulate forwards (KF): given policy functions, push the density forward. Between shock ages: KF equation (sparse matrix, fast). At shock ages: apply the discrete redistribution cascade

Step 4 — Check market clearing: compute labour supply, housing demand from g_t; compare to labour demand, housing supply. Update prices and return to Step 1

Computation times (M1 Ultra, MATLAB)

San Francisco (55 locations, 3,025 pairs): steady state = **9 seconds**. New York (183 locations, 33,489 pairs): steady state = **82 seconds**; transitions ≈ **2.3 hours**. Without shock ages, the New York problem would be infeasible.

Validation: Non-Targeted Moments

Figure 2: Non-targeted moments

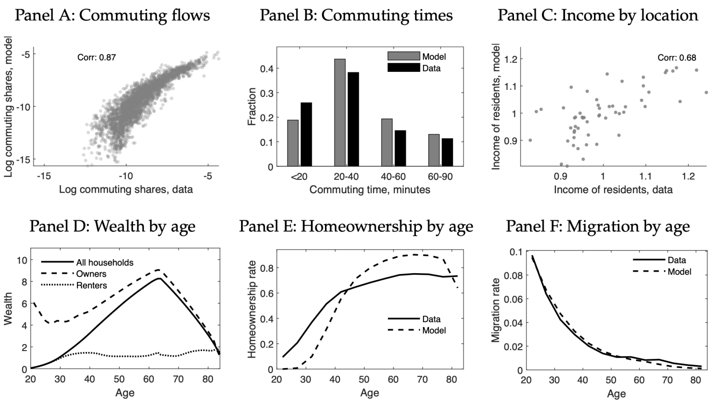

The model reproduces (without targeting): commuting time distribution, spatial income gradient, hump-shaped homeownership profile, declining migration rate over the life cycle. The homeownership profile validates the tenure-choice mechanism before any counterfactual.

The Three Experiments

1. Productivity shock
Tech-sector TFP boom (2009–2015) in a subset of Bay Area locations. Calibrated using the shift-share (Bartik) instrument (Eq. 4.1). Natural experiment in the data.

Question: who gains and who loses from a local productivity surge?

2. Upzoning
Increase construction productivity Z^h_{Sit} in locations with below-median construction productivity. Reduces the cost of supplying new housing.

Question: does welfare-improving housing supply expansion face political opposition?

3. High-Speed Rail
Introduce four stations of the planned California HSR, reducing commuting times to those stations.

Question: do infrastructure improvements generate winners and losers through the housing market?

The common thread

All three experiments have positive aggregate welfare effects. All three produce a minority of losers — and the identity of the losers depends critically on homeownership status and location.

Productivity Shock: Aggregate Effects and the Transition

Key finding: ignoring the transition *underestimates* welfare gains by more than a factor of two.

	With transition	Steady-state only
Aggregate welfare gain	**1.37%**	0.59%
Long-run population change	+6.31%	

Why the gap is so large:

- House prices *jump on impact* when the shock is announced (forward-looking agents anticipate future wage gains)
- Current homeowners capture the price appreciation immediately — a gain that does not appear in the steady-state comparison
- Sluggish migration and durable housing mean the transition itself generates welfare gains (above and beyond the new steady state)

A static model calibrated to the Bay Area productivity shock systematically underestimates aggregate welfare gains by a factor of two. Every policy simulation using a static QSM has this bias.

Transition Dynamics: The Path Matters

Figure 6: Transitional effects of the productivity shock

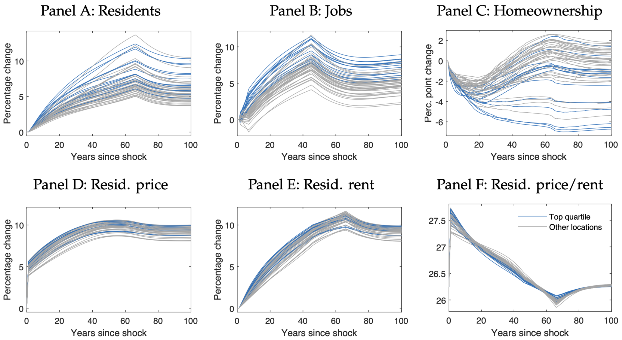

Non-monotone transitions

Some locations *lose jobs in the first few years* after the shock before gaining them in the long run. The static model misses this entirely. A policymaker evaluating the shock at the 5-year horizon would reach different conclusions from one evaluating at the 20-year horizon.

How Long Does Adjustment Take?

Transition half-lives (time for a median location to move halfway to the new steady state):

	Residents	Jobs
Productivity shock	**19 years**	**6 years**
High-Speed Rail	**28 years**	**14 years**

Why jobs adjust faster than residents:

- Employment responds through commuting (the fast margin — workers can commute to new jobs without moving)
- Residential reallocation requires house purchases, sales, and moving costs (the slow margin — fully dynamic, 19-year half-life)
- This is the DUE counterpart of MRR's key result: commuting and migration are different adjustment margins operating at different speeds

Connection to Lecture 6 (MRR)

MRR showed cross-sectional heterogeneity in employment elasticities. DUE reveals the time dimension: even in a county with high employment elasticity, the residential adjustment takes nearly two decades.

The Homeownership Wedge: Who Wins and Why

Comparing three models (Table 3): main model vs. no-homeownership (NHO) vs. hand-to-mouth (HTM).

	Main	NHO	HTM
Mean welfare gain	1.37%	0.97%	0.26%
Welfare dispersion (p95−p5)	3.82 ppt	2.00 ppt	1.59 ppt
Share harmed	**3.9%**	**25.0%**	**50.7%**

Why the models diverge so sharply:

- In the main model, owner-occupied housing is a large, leveraged, spatially undiversified asset — house price jumps generate gains concentrated on owners
- NHO exposes households to real estate only via REITs; the jump in individual home values is absent; 25% lose because they cannot hedge rising housing costs
- HTM cannot insure against future costs at all — 51% lose even though the aggregate shock is positive

Tenure status at the moment of the shock is the primary determinant of individual welfare. A model without homeownership does not just get the numbers wrong — it gets the distribution wrong.

Distributional Welfare Effects: Age and Tenure

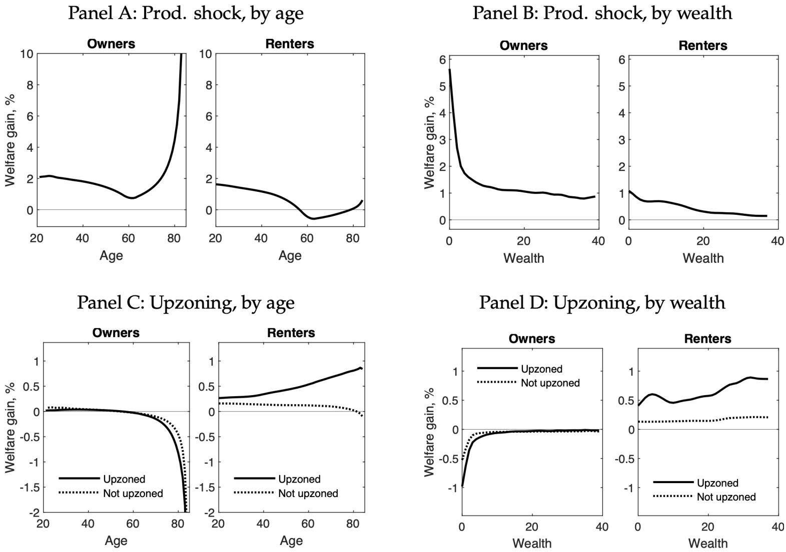

Upzoning and the NIMBY Coalition

Aggregate results (Table 2): upzoning raises average welfare by **0.12%** and population by 0.64% in the long run. Rents fall. Construction increases.
But: in *every single location*, the majority of homeowners experience welfare losses.

	Main	NHO
Mean welfare gain	0.12%	0.19%
Share of all households harmed	**40.6%**	6.8%

Why homeowners lose:

- Upzoning increases housing supply $\Rightarrow$ rents and prices fall
- Homeowners hold an illiquid primary asset whose value declines
- Capital loss on the leveraged housing position $>$ gain from lower future rents
- Old HO most exposed: large house, short remaining horizon, close to selling

The rational NIMBY

This opposition is not irrational or ill-informed. Homeowners are correctly calculating their welfare under the policy. A majority of the most politically active voters in local planning decisions face genuine losses from upzoning.

Closing the Loop from Lecture 1

Hsieh & Moretti (2019)
The aggregate cost of housing constraints:

- Removing zoning constraints in NY, SF, SJ would raise aggregate US wages by 2–9%
- Constraint = aggregate misallocation

The puzzle left open: if constraints are so costly, why do they persist? Who opposes reform?

DUE (Greaney et al. 2025)
Why constraints persist:

- Upzoning harms the majority of homeowners in every location
- Homeowners are concentrated in the voting population at local planning meetings
- Opposition is rational: it is the correct response to welfare losses from illiquid asset devaluation

The course arc

Lecture 1 documented the aggregate cost. DUE explains why the households who bear the cost of removing constraints are the same households with the most political salience in local planning decisions. The welfare-improving policy faces a structural political economy barrier rooted in the distribution of housing wealth.

Three Things CDP Establishes

Before the model, the three results that motivate it:

1. **Aggregate gains $\neq$ distributional costs**: the US gains from the China shock, but mobility frictions mean manufacturing workers in exposed states bear concentrated losses while service workers elsewhere gain
2. **Frictions convert a trade shock into a prolonged regional crisis**: the persistence ADH documented is evidence that frictions are large; CDP estimates the full adjustment path — the national-scale counterpart of DUE's 19-year half-transition
3. **Static welfare comparisons mislead**: comparing steady states ignores transition costs ($\approx$ **4.7%** of long-run welfare gains) and misattributes who gains

Each structural element that follows — forward-looking migration, trade gravity, input-output linkages, dynamic hat algebra — is there to deliver one of these three conclusions at national scale with full GE accounting. Framework: 38 countries / 22 sectors / 50 US states.

Scaling the Question: From City to Nation

DUE
Within one city:

- Shock: tech productivity boom
- Scale: 55 locations, 3,025 pairs
- Key asset: owner-occupied housing
- Initial sorting: tenure status

CDP
Across the nation:

- Shock: China trade shock (2000–07)
- Scale: 38 countries, 22 sectors, 50 states
- Key friction: inter-sector/inter-state mobility
- Initial sorting: sector $\times$ state location

The China trade shock

US imports from China more than doubled, 2000–2007. Manufacturing employment fell sharply in exposed regions. Autor, Dorn & Hanson (2013) documented the reduced-form regional losses. CDP asks what ADH cannot: how long did adjustment take, and what did full GE distribute?

CDP: Model Scale and Key Equations

Scale: $N = 38$ countries, $J = 22$ sectors, $R = 50$ US states. A US labor market $=$ sector j in state n ($\approx$ 1,150 markets including non-employment).

Household utility (Eq. 1): Cobb-Douglas over final goods, log utility

$$C_t^{nj} = \prod_{k=1}^{J} (c_t^{njk})^{\alpha_k}, \qquad U(C_t^{nj}) = \log C_t^{nj}$$

Household value function (Eq. 2): forward-looking, discrete-choice migration

$$V_t^{nj} = U(C_t^{nj}) + \nu \log\left(\sum_i \sum_k \exp\left(\frac{\beta V_{t+1}^{ik} - T^{nj,ik}}{\nu}\right)\right)$$

$\beta = 0.99$ (quarterly); $\nu = 5.34$; $T^{nj,ik} =$ cost of relocating from nj to ik.

Migration shares (Eq. 3):

$$\mu_t^{nj,ik} = \frac{\exp\big((\beta V_{t+1}^{ik} - T^{nj,ik})/\nu\big)}{\sum_m \sum_h \exp\big((\beta V_{t+1}^{mh} - T^{nj,mh})/\nu\big)}$$

Trade Gravity and Dynamic Hat Algebra

Trade shares (Eq. 7 — EK gravity):

$$\pi_t^{nj,ij} = \frac{(x_t^{ij}\ \kappa_t^{nj,ij})^{-\theta^j}\ (A_t^{ij})^{\theta^j/\gamma^{nj}}}{\sum_m (x_t^{mj}\ \kappa_t^{nj,mj})^{-\theta^j}\ (A_t^{mj})^{\theta^j/\gamma^{nj}}}$$

θ^j = trade elasticity by sector; $\kappa_t^{nj,ij}$ = bilateral trade cost (i to n); x_t^{ij} = unit cost bundle; A_t^{ij} = sectoral-regional TFP; γ^{nj} = value added share.

Dynamic hat algebra (Propositions 1–3): express equilibrium conditions in time differences. All unobserved fundamentals (TFP levels, trade costs, mobility costs) difference out.

What you need

- Observed initial allocations: employment shares, migration flows, trade shares
- Three parameters: θ^j (Caliendo & Parro 2015), $1/\nu = 0.2$ (ACM 2010), $\beta = 0.99$
- The counterfactual: calibrated change in Chinese TFP by sector

This is the dynamic analog of Dekle-Eaton-Kortum (2008) exact hat algebra —

CDP: Production and Inter-Sectoral Linkages

Intermediate goods: firms in sector j, region n produce differentiated varieties using labour l, structures h, and materials M from *all* sectors:

$$q_t^{nj} = z^{nj}\left(A_t^{nj}\right)(h_t^{nj})^{\xi^n}(l_t^{nj})^{1-\xi^n}\prod_{k=1}^{J}(M_t^{nj,nk})^{\gamma^{nj,nk}}$$

γ^{nj} = value added share; $\gamma^{nj,nk}$ = input-output coefficient (materials from sector k used in j). Constant returns; structures in fixed supply.

Trade: varieties traded across regions at iceberg cost $\kappa_t^{nj,ij}$. Competitive firms $\Rightarrow$ price = minimum unit cost across origins (EK). Trade shares $\pi_t^{nj,ij}$ follow gravity.

Why this matters: a productivity shock in Chinese electronics does not just hit US electronics workers. Through input-output linkages, it lowers the cost of materials for US machinery, construction, and other downstream sectors. Through trade, it changes the price index in every region. GE effects propagate across *both* sectors and regions — which is why a 22-sector model is needed.

Market clearing: labour $L_t^{nj} = \gamma^{nj}(1-\xi^n)\sum_i \pi_t^{ij,nj} X_t^{ij}/w_t^{nj}$; goods expenditure X_t^{nj} links demand from all sectors/regions via IO and trade shares.

CDP: Three Equilibrium Concepts

The state of the economy at time t is the distribution of labour $L_t = \{L_t^{nj}\}$ across markets. CDP defines three nested equilibrium concepts:

1. Temporary equilibrium (Def. 1): given L_t and fundamentals $(\Theta_t, \bar{\Theta})$, solve for wages w_t and allocations that clear all goods and labour markets *at time t*. This is a static multi-country, multi-sector trade model (Caliendo & Parro 2015).

2. Sequential competitive equilibrium (Def. 2): given initial L_0 and a path of fundamentals $\{\Theta_t\}$, a sequence $\{L_t, \mu_t, V_t, w_t\}$ such that the temporary equilibrium holds at every t *and* labour evolves via forward-looking migration: $L_{t+1}^{nj} = \sum_{i,k} \mu_t^{ik,nj} L_t^{ik}$.

3. Stationary equilibrium (Def. 3): a sequential equilibrium in which all aggregate variables are constant. Households may still move, but inflows and outflows balance everywhere.

The layered structure

Temporary equilibrium is a building block: at each t, given labour supply, solve for prices. The sequential equilibrium links periods through forward-looking migration. Dynamic hat algebra (Propositions 1–3) solves the sequential equilibrium in *time differences*, avoiding estimation of any fundamental level.

China Shock: Identification

Instrument: ADH-style — predicted change in US imports from China using actual change in imports of other high-income countries as IV:

$$\Delta M^{\text{USA},j} = a_1 + a_2 \, \Delta M^{\text{other},j} + u_j, \qquad \hat{a}_2 = 1.386 \; (R^2 = 0.99)$$

Calibrated Chinese TFP shock (2000–2007):

- Average manufacturing TFP growth ≈ **11%**
- Computer & electronics: ≈45% of total US import increase from China
- Textiles: ≈15%; metals and machinery: ≈7% each

Key data moment motivating large frictions (Table I, quarterly):

Transition type	Median quarterly rate
Changing sector (same state)	5.44%
Changing state (same sector)	0.42%
Staying in same market	93.9%

Employment Effects: Aggregate and Sectoral

Main result (calibrated counterfactual, 2000–2007):

- China shock reduced manufacturing employment share by **0.36 percentage points**
- $\approx$ **0.55 million** manufacturing jobs — about 16% of observed decline
- Construction gained $\approx$50,000 jobs (cheaper imported intermediates lower costs)
- Long-run non-employment share: -0.22 ppt (modest improvement)

Sectoral concentration: computer & electronics $\approx$25% of total manufacturing decline; furniture, textiles, metals, machinery each 10–15%.

The shock is sectorally concentrated, not manufacturing-wide. Workers in unexposed manufacturing sectors in the same state are collateral damage through GE linkages, not through direct trade exposure.

Dynamic Employment Path

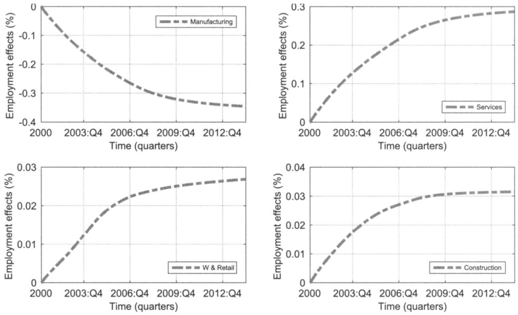

What Figure 1 shows

- Manufacturing employment falls immediately on impact and continues declining
- Inter-sector and inter-state mobility is costly — $\nu = 5.34$ implies slow reallocation
- Services rise gradually; the full adjustment takes longer than the observed 2000–2007 window

Welfare Effects: Aggregate and Distributional

US aggregate welfare: $+\mathbf{0.2}\%$ from the China shock.

Distribution across 1,150 US labor markets:

- Range: from -0.8% to $+1.0\%$
- On impact: real wages fall in $\approx$ **47%** of labor markets
- Long run: only $\approx$ **4%** of labor markets have welfare losses
- All 50 US states gain in the long run (Michigan: $+0.12\%$; Vermont: $+0.22\%$)

The aggregate $+0.2\%$ masks large cross-sectional dispersion. Manufacturing workers in exposed states bear losses for many periods while non-manufacturing workers elsewhere gain immediately. Reporting only the aggregate is misleading about who bears the cost.

Welfare Distribution Across Labor Markets

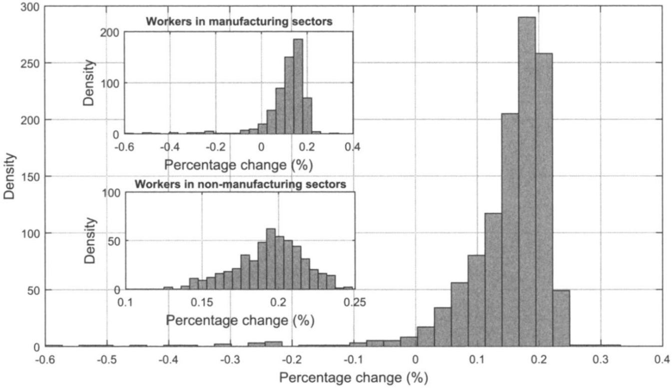

The key pattern

Manufacturing workers in exposed sectors and states bear concentrated short-run losses; service-sector workers in the same states gain from cheaper goods. This is the aggregate-scale counterpart of DUE's homeownership wedge: initial sectoral location determines welfare outcome, just as tenure status does in DUE.

Transition Costs and Mobility Frictions

Transition costs: $\approx$ **4.7%** reduction in the value of long-run welfare gains. Workers pay this fraction of steady-state gains to avoid the transition path.

Migration elasticity: $1/\nu = 0.2$ ($\nu = 5.34$, quarterly). Estimated via GMM following Artuç, Chaudhuri & McLaren (2010); implied annual $\nu = 2.02$, very close to ACM's $\nu = 1.88$.

Robustness: allowing Poisson re-optimization shocks with arrival rate $p \in \{0, 0.1, 0.2, 0.3\}$ barely changes manufacturing employment effects (range: -0.296 to -0.361 ppt).

Mobility frictions are the engine converting a one-time trade shock into a prolonged regional crisis. Without frictions, adjustment is immediate. With $\nu = 5.34$, the welfare cost of being in the wrong market accumulates over many periods.

DUE and CDP: Two Scales of the Same Friction

DUE (city scale)

- Shock: tech productivity boom in SF Bay Area
- Friction: illiquid owner-occupied housing
- Mechanism: house price jump distributes gains to owners; future buyers and renters face higher costs
- Consequence: 40.6% of households harmed by welfare-improving upzoning
- Half-life: 19 years (residents)

CDP (national scale)

- Shock: China trade shock (2000–2007)
- Friction: labor mobility across sector-state markets
- Mechanism: manufacturing workers in exposed states cannot reallocate quickly
- Consequence: 47% of labor markets see real wage drops on impact despite positive aggregate
- Transition cost: 4.7% of long-run welfare

The shared lesson

In both papers: aggregate welfare gains are real, but frictions mean the gains and losses fall on *different people*. The static comparison (comparing steady states) misses the distributional costs of transition. Dynamic frameworks are not a refinement — they change the fundamental welfare calculation.

Initial Sorting as the Central Determinant

Both papers take the initial distribution of households as given:

- DUE starts from observed homeownership and location patterns in Bay Area
- CDP starts from observed sectoral and regional employment distributions across 50 states

Both papers show: initial sorting determines who wins and who loses.

- DUE: tenure status at the moment of the shock determines the welfare outcome; the average welfare gain is 1.37% in the main model vs. 0.97% without homeownership — a 0.40 ppt gap attributable entirely to the asset price channel that homeowners experience
- CDP: sectoral and regional sorting at the start of the China shock determines transition costs; working in manufacturing in Michigan vs. services in Vermont determines whether your welfare falls on impact

The open question

If sorting drives distribution, what produces the equilibrium sorting of households across locations, tenure statuses, and sectors? Neither DUE nor CDP addresses this. The equilibrium sorting is an input, not a result. That is what Lecture 10 opens.

Bridge to Lecture 10

What DUE and CDP jointly establish:

- Transition dynamics matter for welfare measurement — static comparisons systematically bias estimated gains and misdistribute them across agents
- Homeownership is a structural parameter of spatial equilibrium, not a side feature
- Mobility frictions are heterogeneous and partly driven by asset illiquidity

What remains to be asked:

- DUE and CDP both start from an *observed initial equilibrium*. What produced that equilibrium? In particular: why are homeownership rates high in exactly the places most exposed to adverse shocks?
- If young workers rent and old workers own, and if the young are more mobile, then the housing market *sorts* workers by mobility into tenure categories — and this sorting is intergenerational
- **Lecture 10** (Guaitoli, Pancrazi & Raimondo 2026; Fajgelbaum & Gaubert 2020): what determines the sorting of households into locations and tenure statuses, and what normative framework governs optimal spatial policy?

Summary: Two Papers, One Arc

DUE
Greaney, Parkhomenko & Van Nieuwerburgh (2025)

Key features: many locations, commuting, forward-looking agents, life cycle, tenure choice, illiquid housing, endogenous supply

Key result: productivity shock welfare gain is 1.37% (with transition) vs. 0.59% (steady-state) — factor of 2; upzoning harms majority of homeowners in every location

Core takeaway: tenure status at shock time is the primary welfare determinant; NIMBY opposition is rational

CDP
Caliendo, Dvorkin & Parro (2019)

Key features: 38 countries, 22 sectors, 50 US states; EK trade, forward-looking migration, dynamic hat algebra

Key result: China shock reduces US manufacturing by 0.55M jobs (0.36 ppt); aggregate welfare $+0.2\%$ but 47% of labor markets see wage drops on impact; transition costs $\approx 4.7\%$

Core takeaway: frictions make trade shocks prolonged regional crises; aggregate gains mask concentrated distributional losses

References

- Greaney, Parkhomenko & Van Nieuwerburgh (2025). "Dynamic Urban Economics." Working paper, December 2025.
- Caliendo, Dvorkin & Parro (2019). "Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock." *Econometrica* 87(3): 741–835.
- Ahlfeldt, Redding, Sturm & Wolf (2015). "The Economics of Density: Evidence from the Berlin Wall." *Econometrica* 83(6): 2127–2189.
- Autor, Dorn & Hanson (2013). "The China Syndrome: Local Labor Market Effects of Import Competition." *American Economic Review* 103(6): 2121–2168.
- Artuç, Chaudhuri & McLaren (2010). "Trade Shocks and Labor Adjustment." *American Economic Review* 100(3): 1008–1045.
- Monte, Redding & Rossi-Hansberg (2018). "Commuting, Migration, and Local Employment Elasticities." *AER* 108(12): 3855–3890.
- Glaeser & Gyourko (2005). "Urban Decline and Durable Housing." *Journal of Political Economy* 113(2): 345–375.
- Achdou, Han, Lasry, Lions & Moll (2022). "Income and Wealth Distribution in Macroeconomics." *Review of Economic Studies* 89(1): 45–86.
- Hsieh & Moretti (2019). "Housing Constraints and Spatial Misallocation." *AEJ: Macro* 11(2): 1–39.