Macro-Spatial Economics

Lecture 8: Housing, Regional Shocks, and Aggregate Fluctuations

Roberto Pancrazi

University of Warwick

Where We Are in the Course

The Arc So Far

1. **L5**: Welfare and incidence — who gains and who loses from local shocks (DR; HM 2024)
2. **L6**: Adjustment margins — commuting, migration, housing frictions (MRR; Notowidigdo)

Today's Question

We know *that* local shocks matter and *why* workers do not always move. Now: **what do regional shocks and regional data tell us about the causes of the Great Recession and aggregate business cycles?**

Today's Papers

- **Part I**: Mian & Sufi (2014, *Econometrica*) — the regional evidence and its limits
- **Part II**: Beraja, Hurst & Ospina (2019, *Econometrica*) — the structural fix

The Intellectual Thread

Both papers ask the same question: **what caused the 2007–2009 employment collapse?**

Mian & Sufi (2014)

Strategy: exploit cross-county variation in housing net worth shocks

Answer: housing net worth channel explains non-tradable job losses

Implicit conclusion: demand shocks + some wage rigidity

BHO (2019): The Critique

Problem: back-of-the-envelope extrapolation from regional elasticities *risks missing channels and shocks important at the aggregate level but not the regional level*

Fix: embed regional variation in a structural model

This lesson's arc

You must see the regional approach *first* to understand what BHO are critiquing and what their methodology adds. The lesson is about the *limits* of regional identification, not just its power.

The Great Recession: The Facts to Explain

Stylized Facts, 2007–2009

- Employment-to-population ratio fell from 63% to 58% — **8.6 million jobs lost**
- Largest employment collapse since the Great Depression
- Simultaneously: house prices fell roughly 30% nationally; household net worth fell $13 trillion
- Consumption fell sharply before the financial panic of autumn 2008

The central observation: the collapse in employment coincided almost perfectly with the collapse in housing net worth — both geographically and temporally.

The Mian & Sufi Research Strategy

The housing bust was *unequal* across U.S. counties: some counties lost 45% of net worth, others experienced slight gains. This geographic variation is an identification opportunity.

The Housing Net Worth Channel: Mechanism

Definition (Mian & Sufi 2014)

The **housing net worth channel**: a decline in housing net worth suppresses consumer demand, either through a direct wealth effect or through tighter borrowing constraints (collateral channel), which then reduces local employment.

The testable prediction: if the channel operates through *local demand*, then:

- **Non-tradable** employment (retail, restaurants) should fall more in counties with large HNW declines — it relies on local demand
- **Tradable** employment (manufacturing, exports) should be *unaffected* — it sells to national/global markets

Why the Tradable/Non-Tradable Distinction Is Crucial

If both tradable and non-tradable employment fell equally, we could not distinguish demand from supply shocks (e.g. a credit tightening or productivity shock would hit all sectors). The asymmetric prediction is what pins down the channel.

M&S (2014): A Simple Labor Market Model

M&S formalise the mechanism with a two-sector model (Section 3) to clarify what adjustment *would* look like if wages and workers were flexible.

Setup: Tradable (T) and Non-Tradable (NT) Sectors

Workers are mobile across sectors within a county (but not across counties). A housing net worth shock reduces local demand for NT goods.

Case 1 — Flexible wages: NT labor demand falls $\Rightarrow$ NT wages fall $\Rightarrow$ workers shift to T sector $\Rightarrow$ T employment *rises*, NT employment falls, total employment nearly unchanged.

M&S (2014): A Simple Labor Market Model

Case 2 — Sticky wages: wages cannot fall $\Rightarrow$ NT workers are laid off and *remain unemployed* $\Rightarrow$ T employment unaffected $\Rightarrow$ total employment falls sharply.

The Model's Prediction

If adjustment works: tradable employment should *expand* in hard-hit counties (T absorbs displaced NT workers).
If adjustment fails: tradable employment is flat and total employment collapses.
M&S find: tradable employment is flat $\Rightarrow$ consistent with sticky wages *or* immobile workers.

Net Worth: Definition and Measurement

County-Level Net Worth (Mian & Sufi 2014)

$$NW_i^t = S_i^t + B_i^t + H_i^t - D_i^t$$

where S: stocks, B: bonds/deposits, H: housing value, D: debt. All at county level.

Data sources:

- S_i, B_i: IRS Statistics of Income (county-level)
- H_i: 2000 Decennial Census $\times$ CoreLogic ZIP-code house price index
- D_i: Equifax Predictive Services (total household borrowing by county)

Housing net worth shock:

$$\Delta HNW_i = \frac{\Delta \log p_i \cdot H_i^{2006}}{NW_i^{2006}}$$

Average ΔHNW 2006–2009: -9.5% (weighted mean), with **large dispersion** across counties — the key source of identification.

Industry Classification: Tradable vs. Non-Tradable

Two independent methods (Mian & Sufi 2014):

Method 1 — Trade-flow based:

- Tradable: imports + exports $\geq$ \$10,000/worker or total trade $\geq$ \$500M
- Non-tradable: retail sector + restaurants
- Construction: separate category, excluded from both

Method 2 — Geographic concentration:

- Tradable industries need scale and specialisation $\Rightarrow$ geographically concentrated (Herfindahl index 0.018)
- Non-tradable industries serve local demand everywhere $\Rightarrow$ geographically dispersed (Herfindahl index 0.004)

Robustness

Both classification methods deliver the same empirical results. The concentration-based method also captures activities like amusement parks and internet service providers as tradable — which would be missed by a pure trade-flow approach.

Summary Statistics: The Cross-County Variation

Variable	Mean
Housing net worth shock, 2006–2009	
ΔHNW_i (population-weighted)	-9.5%
Bottom decile of counties	-45%
Top decile of counties	+slight gain
Employment shares (2007)	
Non-tradable	20%
Tradable	11%
Construction	11%
Other (government, etc.)	59%

Key message: the cross-county dispersion in ΔHNW is enormous — this is what makes identification possible. A county at the bottom decile experienced a balance-sheet shock roughly **five times** larger than the average.

Source

County Business Patterns (CBP) for employment by 4-digit NAICS; Equifax + CoreLogic for net worth; ACS for wages. Sample: all U.S. counties, 2007–2009.

Identification: The Housing Supply Elasticity Instrument

Problem: counties with large housing net worth declines may also differ in labour demand, local institutions, or industry mix $\Rightarrow$ OLS is biased.

Instrument: Saiz (2010) Housing Supply Elasticity

Areas with inelastic housing supply (mountains, coastlines, regulations) experienced larger boom–bust cycles in house prices. The *topographic* constraint on supply is plausibly exogenous to local labour demand shocks.

First stage: low supply elasticity $\Rightarrow$ large ΔHNW 2006–2009.
Exclusion restriction: topography affects employment only through house prices.

Connection to Lecture 6

The Saiz instrument also appears in Notowidigdo (2020, Lecture 6): housing supply inelasticity traps workers in declining areas. Here it identifies the size of the balance sheet shock. **Same instrument, two related stories about housing frictions.**

Main Result: Non-Tradable Employment

Baseline Regression (Mian & Sufi 2014)

$$\Delta \log E_i^{NT} = \alpha + \beta \cdot \Delta HNW_i + \mathbf{X}_i'\gamma + \varepsilon_i$$

E_i^{NT}: non-tradable employment in county i, 2007–2009. IV: Saiz elasticity.

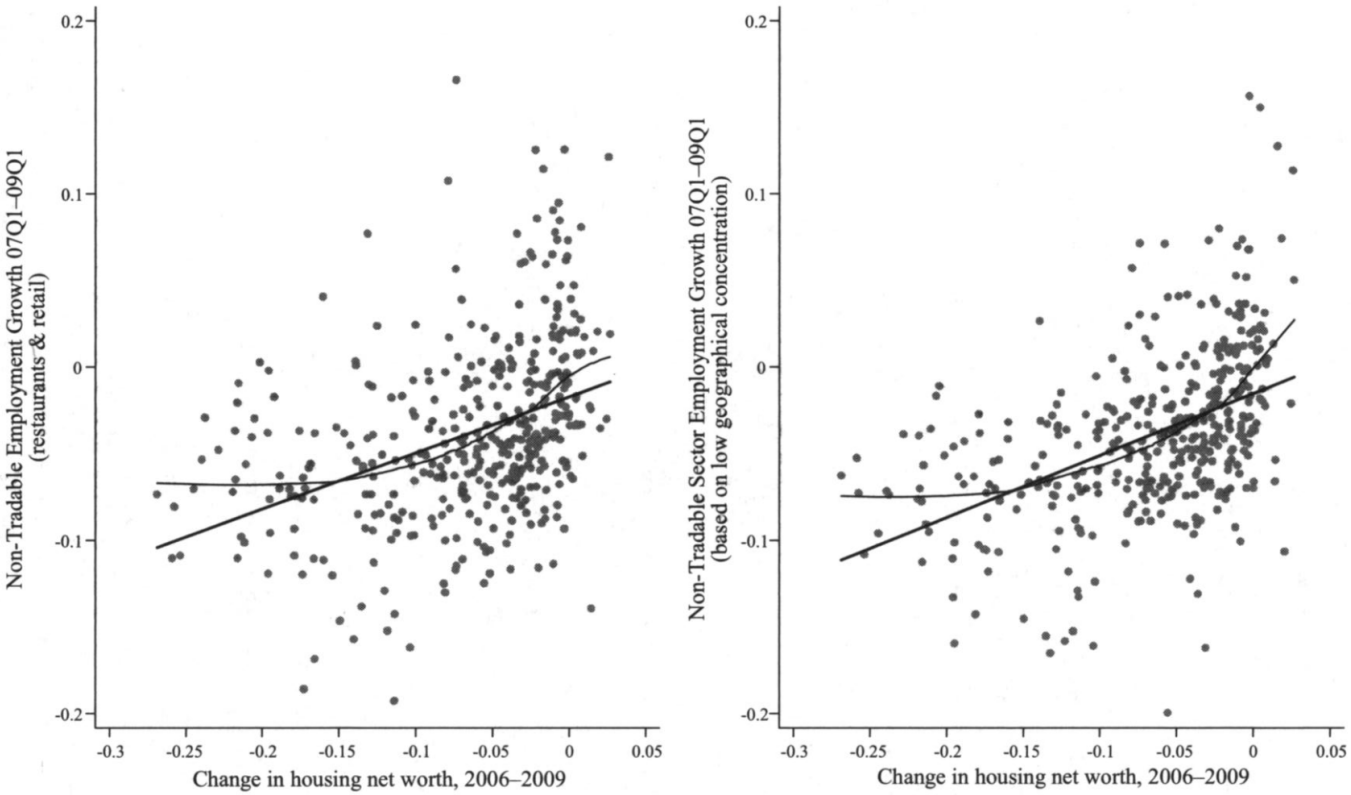

Main Result: Quantitative Findings

Baseline Regression (Mian & Sufi 2014)

$$\Delta \log E_i^{NT} = \alpha + \beta \cdot \Delta HNW_i + \mathbf{X}_i'\gamma + \varepsilon_i$$

E_i^{NT}: non-tradable employment in county i, 2007–2009. IV: Saiz elasticity.

Key finding: a **10 percentage point** decline in housing net worth is associated with a **3.7 percentage point** decline in non-tradable employment.

Result is **not** driven by:

- Industry-specific supply shocks (controlled by 23 two-digit industry shares)
- Construction exposure (direct controls + IV)
- Credit tightening to firms (effect stronger for *large* establishments, less credit-constrained)

The Placebo: Tradable Employment

Tradable Employment Regression

$$\Delta \log E_i^T = \alpha + \beta \cdot \Delta HNW_i + \mathbf{X}_i'\gamma + \varepsilon_i$$

Finding: *zero* correlation between ΔHNW and tradable employment growth.

This is the **key falsification test**. A demand channel operating through local spending should:

- Hit non-tradable employment: ✓ (3.7 pp per 10 pp HNW decline)
- Leave tradable employment unaffected: ✓ (near-zero coefficient)

What This Rules Out

If the employment decline were due to a credit supply shock (banks pulling back on lending) or a technology shock, we would expect tradable employment to fall too. The tradable placebo discriminates the demand channel from supply-side alternatives.

Labor Market Adjustment: Wages and Migration

Mian & Sufi (2014) also test whether adjustment mechanisms are operating.

Wage Adjustment

Counties most affected by HNW decline show **little nominal wage response**. Local wages appear *sticky*: they do not fall more in harder-hit counties.

Implication: wage rigidity may be preventing the labour market from clearing.

Migration Adjustment

Little evidence of out-migration from counties with the largest HNW declines.

This echoes Notowidigdo (2020, Lecture 6): housing costs and transfers impede geographic reallocation.

The Back-of-the-Envelope Conclusion

Housing net worth shock $\rightarrow$ demand collapse $\rightarrow$ non-tradable employment falls. Wages and workers don't adjust. Scaling up: the aggregate employment drop looks like a *demand-driven recession with sticky wages*.
But is this a valid extrapolation from regional to aggregate?

The Limits of the Regional Approach

Mian & Sufi (2014) use cross-sectional variation to identify a *local* elasticity, then implicitly scale it to the aggregate.

The Back-of-the-Envelope Logic

1. Identify: housing shock $\rightarrow$ 3.7 pp non-tradable employment decline per 10 pp HNW decline
2. Scale: aggregate HNW fell 9.5% on average
3. Conclude: most of the aggregate employment decline is explained by the housing demand channel + wage rigidity

The Limits: BHO's Critique (Preview)

Two Errors in the Extrapolation

This risks **two errors**:

1. **Structural**: the same type of shock may produce different regional vs. aggregate responses (monetary policy responds to aggregate, not regional, shocks)
2. **Shock composition**: aggregate-only shocks (labor supply) are *differenced out* in regional regressions but were important for the aggregate recession

The conclusion about wage rigidity is particularly fragile — wages look more flexible in regional data than in aggregate time series.

The Aggregate Wage Puzzle

Two Facts in Tension

Regional (cross-state, 2007–2010): A 1% decline in state employment was associated with a **0.72%** decline in nominal wages and **0.64%** in real wages.

Aggregate (time series, 2007–2010): Aggregate employment fell 7.7%. Yet nominal wages barely moved — aggregate nominal wage elasticity ≈ 0.51–0.54; real wage elasticity ≈ 0.34–0.37.

The Puzzle

Wages appear *twice as flexible* in regional data as in aggregate time series. Standard NK models calibrated to aggregate data imply only $\approx 50\%$ annual wage reset, consistent with the aggregate pattern. But if wages are that sticky, why is the regional elasticity so high?

Regional vs. Aggregate: Wage elasticity

	Nominal	Real
Cross-state, 2007–2010 (BHO Table I)		
Regional	0.72 (se=0.14)	0.64 (se=0.16)
Aggregate time series, 2007–2010 (BHO Table II)		
CPS	0.51	0.34
ACS	0.54	0.37

Regional elasticity is roughly twice the aggregate elasticity.

What Makes This Surprising

The aggregate is the sum of the regions — yet wages look twice as flexible in regional data. Not measurement error: ACS/CPS aggregate correlation is 0.98; same demographic adjustment in both.

The Inference Risk

Regional elasticity (0.72) $\Rightarrow$ wages flexible, demand shocks overstated. Aggregate elasticity (0.51) $\Rightarrow$ sticky wages, large demand shocks. **Both cannot be right simultaneously.**

Why Regional $\neq$ Aggregate: Two Reasons

BHO identify two reasons why wage and employment co-movement differs between regional and aggregate data:

Reason 1 — Different structural responses to the same shock type:

- For a household demand shock: regional wage elasticity can be *smaller* than aggregate because of cross-region trade in intermediates and borrowing/lending
- But aggregate elasticity may be *smaller* because the interest rate endogenously responds to aggregate shocks (Taylor rule) but not to regional shocks

Reason 2 — Different shock mix:

- Regional regressions *difference out* shocks with no cross-regional variation — including aggregate labour supply shocks
- Aggregate labour supply shocks push wages and employment in *opposite* directions $\Rightarrow$ wages look sticky in aggregate but not regionally

Implication

You cannot simply scale a regional elasticity to aggregate conclusions. The structural parameters may be shared, but the observed elasticities embed both the structural response *and* the shock distribution.

The Observational Equivalence Problem

Story A: Sticky wages

ξ_w high (50% reset/yr): wages rigid

Large demand shock $\Rightarrow$ employment collapses, wages flat

Aggregate: emp $\downarrow$ 7.7%, wages flat ✓

Story B: Flexible wages

ξ_w low (76% reset/yr): wages respond

Demand shock $\Rightarrow$ emp and wages both fall + Labour supply shock $\Rightarrow$ wages pushed back up

Net: emp $\downarrow$ 7.7%, wages flat ✓

Why the Aggregate Cannot Distinguish Them

Both fit the aggregate data equally well (similar log-marginal likelihood in BHO).

Different policy implications: Story A $\Rightarrow$ demand stimulus effective; Story B $\Rightarrow$ stimulus less effective, labour supply is the margin.

The Model: A Monetary Union of Islands

Setup (BHO 2019, Section 4)

- Continuum of **islands** k, each with households and firms in two sectors
- **Intermediate goods**: traded across islands (law of one price)
- **Final consumption good**: non-tradable (island-specific)
- **Labor**: mobile across sectors within island, *immobile* across islands
- **Assets**: tradable one-period nominal bond; non-tradable physical capital
- **Monetary policy**: common interest rate R_t set by a union-level Taylor rule

Why No Migration?

Labor immobility is a simplifying assumption to isolate the Phillips curve mechanism. MRR (2018, Lecture 6) showed migration matters for long-run adjustment — BHO abstract from this to sharpen identification of the wage-stickiness parameter at business cycle frequencies.

Households and Wage Setting

Households on island k maximize:

$$E_0 \sum_{t=0}^{\infty} \beta^t \hat{b}_{kt} \left[\log(C_{kt} - hC_{k,t-1}) - \hat{\phi}_{kt} \frac{N_{kt}(j)^{1+\nu}}{1+\nu} \right]$$

h: external habit; ν: inverse Frisch elasticity; $\hat{b}_{kt}$: discount factor shock (household demand); $\hat{\phi}_{kt}$: leisure preference shock (**labour supply shock**).

Calvo Wage Frictions (Erceg, Henderson & Levin 2000)

Each period, fraction $1 - \xi_w$ of households re-optimise their wage. The remaining fraction indexes to lagged inflation with parameter ι_w.

Key parameter: $\xi_w \in (0, 1)$ — the Calvo wage-stickiness parameter.
Fraction of wages adjusting per year: $1 - \xi_w^4$ (quarterly model).
This is the parameter BHO want to identify using regional data.

The Nine Shocks: What Gets Identified Where

BHO allow nine shocks, each with a regional and an aggregate component.

Shock	Regional component?
Household demand ($\hat{b}_{kt}$)	✓
Labour supply ($\hat{\varphi}_{kt}$)	✓*
Investment efficiency ($\hat{\mu}_{kt}$)	✓
Gov't spending ($\hat{\varepsilon}^{g}_{kt}$)	✓
Price markup ($\hat{\lambda}^{p}_{kt}$)	✓
Tradable productivity (A^{x}_{kt})	✓
Non-tradable productivity (A^{y}_{kt})	✓
Aggregate TFP growth (z_t)	×
Monetary policy (η_t)	×

* Has local component, but *aggregate* labour supply is differenced out in cross-regional regressions.

Key Insight

The *aggregate* components of labour supply and monetary policy shocks have **no regional footprint** — differenced out automatically in any cross-regional regression. M&S (2014) were blind to them by construction.

The Key Structural Result: Lemmas 1 and 2

Lemma 1 — Aggregation

The aggregate economy behaves as a single representative island. Aggregate shocks map to aggregate responses using the same structural equations as a single island. The aggregate is not simply the sum of region-specific responses.

Lemma 2 — Small Open Economy

Each individual island behaves as a **small open economy** relative to the aggregate. Crucially, the **Regional New Keynesian Wage Phillips Curve** shares the same structural parameters (κ_w, ξ_w) as the **Aggregate NK Wage Phillips Curve**.

Why This Matters for Identification

Parameter ξ_w (wage stickiness) governs both the regional and aggregate WPC. $\Rightarrow$ We can *identify* ξ_w from regional data (where identification is clean) and *impose* it on the aggregate model (where it is hard to separate from shock composition).

Lemmas 1 & 2: Why Parameters Are Shared Across Scales

Intuition for Lemma 2 (small open economy):
Log-linearise the model around a symmetric steady state. Because all islands share the same parameters and the bond market clears at the union level, island k's first-order conditions in deviation from the aggregate ($\tilde{x}_{kt} = x_{kt} - x_t$) satisfy:

- The island-level Euler equation contains no aggregate monetary policy term (it differences out)
- The island's wage Phillips curve becomes:

$$\tilde{\pi}^w_{kt} = \kappa_w \tilde{n}_{kt} + \beta\, \mathbb{E}_t[\tilde{\pi}^w_{k,t+1}] + (\text{island-specific shocks})$$

where $\kappa_w \equiv \frac{(1-\xi_w)(1-\beta\xi_w)}{\xi_w} \cdot \frac{\lambda_w - 1}{\lambda_w(1+\nu)-1}$ is determined entirely by *preference and technology parameters* — the same as in the aggregate WPC

Lemmas 1 & 2: Why Parameters Are Shared Across Scales

Why Monetary Policy Drops Out Regionally

The Taylor rule sets R_t in response to *aggregate* conditions. From island k's perspective, R_t is exogenous and common — it enters each island's equations identically and therefore cancels when taking deviations from the aggregate. **Regions are like small open economies facing a world interest rate they cannot influence.**

The NK Wage Phillips Curves

Regional WPC (log-linearised, tilde = deviation from aggregate; simplified from BHO eq. 2, dropping habit and indexation terms):

$$\tilde{\pi}^w_{kt} \approx \beta\, \mathbb{E}_t\left[\tilde{\pi}^w_{k,t+1}\right] + \kappa_w\, \nu\, \tilde{n}_{kt} - \kappa_w\, \tilde{w}_{kt} + \tilde{\varphi}_{kt}$$

where $\tilde{\pi}^w_{kt}$ = regional wage inflation (deviation from aggregate), $\tilde{n}_{kt}$ = regional employment deviation, $\tilde{w}_{kt}$ = real wage deviation, $\tilde{\varphi}_{kt}$ = local labour supply shock, and $\kappa_w = \frac{(1-\xi_w)(1-\beta\xi_w)}{\xi_w} \cdot \frac{\lambda_w - 1}{\lambda_w(1+\nu)-1}$.

Aggregate WPC (identical structural parameter κ_w; simplified from BHO eq. 4):

$$\hat{\pi}^w_t \approx \beta\, \mathbb{E}_t[\hat{\pi}^w_{t+1}] + \kappa_w\, \nu\, \hat{n}_t - \kappa_w\, \hat{w}_t + \hat{\varphi}_t + \text{other shocks}$$

The Identification Strategy in One Sentence

The regional WPC is purged of aggregate monetary policy confounds (the Taylor rule affects aggregate π^w but differences out across regions), so it cleanly identifies κ_w via GMM using house price changes and lagged macro variables as instruments for regional labour demand.

Demographically Adjusted Wage Index

Wage Index Construction (BHO eq. 1)

$$\text{Wage}_{kt} = \sum_{g=1}^{24} \text{Share}^{g}_{k,\bar{t}} \cdot \text{Wage}^{g}_{kt} \qquad \text{(BHO eq. 1)}$$

24 groups: 6 age $\times$ 4 education bins. Shares $\text{Share}^{g}_{k,\bar{t}}$ are state-specific, fixed at base year $\bar{t} = 2006$.

Why fix the shares? During recessions, low-wage workers disproportionately lose jobs. A raw average *rises* mechanically (composition bias) — wages look artificially sticky. Fixing demographic shares ensures all wage movements reflect genuine within-group changes.

Data: ACS 2001–2014 (state level); CPS (aggregate). ACS–CPS aggregate correlation: 0.98.

Connection to Lecture 1

HM (2019, L1) fixed demographic weights to avoid composition bias when computing counterfactual wages across cities. Same concern, same solution.

Estimation Strategy: Two-Stage

Stage 1 — Identify κ_w from Regional Data (GMM)

Estimate Regional WPC via GMM using **house price changes** as instruments for the marginal rate of substitution, plus **lagged employment and lagged real GDP per capita** as instruments for future wages.

Identifying assumption: local house price variation proxies for household demand shifts and is orthogonal to local labour supply shocks.
Output: $\hat{\kappa}_w$ and s.e. $\Rightarrow$ prior distribution over ξ_w.

Estimation Strategy: Two-Stage

Stage 2 — Estimate Aggregate Model (Bayesian)

Full medium-scale NK DSGE on aggregate time series. Prior from Stage 1 disciplines ξ_w. Compare: (a) regional prior vs. (b) flat prior (aggregate data only).

Connection to L1, L5, L7

BHO use regional variation to discipline a structural parameter, analogous to how HM (2019, L1) and Hornbeck–Moretti (L5) use regional variation. Their innovation: turning a reduced-form regional estimate of κ_w into a Bayesian prior for the aggregate DSGE.

Result 1: Wages Are More Flexible Than Aggregate Data Suggest

	Aggregate data only	**Regional + aggregate**
Annual wage adjustment rate	$\approx 50\%$	$\approx 76\%$
Calvo parameter ξ_w (quarterly)	high	lower
Log-marginal likelihood	similar	similar

Key finding: wages reset roughly **50% more frequently** when regional data inform the estimate ($\approx$76% vs. $\approx$50% annual adjustment). Yet the fit to aggregate time series is nearly identical.

Why Does This Matter?

Aggregate data alone *cannot distinguish* between high wage stickiness and a particular mix of demand vs. supply shocks. Regional data breaks this degeneracy: the high cross-state wage elasticity (0.72) is inconsistent with ξ_w calibrated to aggregate data alone.

Result 2: Shock Decomposition of the Great Recession

BHO Shock Decomposition

- **2007–2010 (the recession)**: aggregate *demand shocks* (discount rate / household demand) are the primary driver of the employment decline — consistent with the housing demand narrative of Mian & Sufi (2014)
- **2010–2014 (the slow recovery)**: demand shocks reverse; *aggregate labour supply shocks* explain why employment remained depressed well after the recession ended

The Twist

Because wages are actually *flexible* ($\approx$76% annual reset, $\xi_w \approx 0.24$), the model cannot generate persistent employment effects from demand shocks alone. The slow recovery is not wage rigidity — it is negative labour supply shocks that have *no regional footprint* and are therefore invisible to Mian & Sufi's identification.

Labour Supply Shocks and the Slow Recovery (2010–2014)

The Puzzle: Why Did Employment Stay Low?

With flexible wages (≈76% annual reset), demand shocks generate little persistence — once reversed, employment should rebound quickly. Yet employment remained depressed 2010–2014.

The Resolution

Negative aggregate labour supply shocks ($\uparrow \hat{\phi}_t$, preference for leisure):

- Disability rolls, early retirement, skill mismatch — workers valued leisure more
- Push wages *up* and employment *down* simultaneously
- **No regional footprint** $\Rightarrow$ completely differenced out in M&S-style regressions

Punchline

Regional identification sees "continued weak demand." The structural model sees negative labour supply. **Same data, opposite diagnosis, opposite policy prescription.**

Where the Back-of-the-Envelope Goes Wrong

BHO construct an explicit comparison showing what happens if you extrapolate from regional elasticities without a structural model.

Three Approaches Compared

1. **Aggregate data only**: finds high wage stickiness; attributes most of 2007–2014 employment decline to demand shocks + rigid wages
2. **Back-of-the-envelope from regional** (Mian & Sufi style): *overstates* the role of household demand shocks in the sluggish post-recession recovery
3. **BHO methodology** (regional + aggregate): finds flexible wages + large labour supply component in the recovery

The Lesson for Regional Identification

Regionally identified elasticities are **local average treatment effects** for shocks that have regional variation. Aggregate-only shocks (labour supply, monetary policy) are **not identified regionally** — they are differenced out. Scaling up implicitly assumes the aggregate was hit by the same shocks that generated regional variation.

The Methodological Thread: What Regional Data Can and Cannot Do

Paper	Identifies	Cannot identify
HM19 (L1)	City wage gains from TFP	Aggregate welfare; GE effects
MRR (L6)	Employment multipliers; commuting gravity	Aggregate multiplier without model
Notowidigdo (L6)	Incidence by housing tenure	Long-run aggregate house price effects
M&S (2014)	Housing demand channel; NT employment	Aggregate-only shocks; wage rigidity vs. shock mix
BHO (2019)	κ_W: WPC slope, shared across scales	Shock composition; monetary policy

The General Principle

Regional variation is a **laboratory** for structural parameters that are scale-invariant (under assumptions). It is **not** a laboratory for the shock distribution, which can differ entirely between regional and aggregate episodes.

Policy Implications

From M&S 2014: What Channel?

- Housing net worth collapse $\rightarrow$ local demand $\rightarrow$ non-tradable employment
- Standard labour market adjustments (wage cuts, migration) did not operate
- $\Rightarrow$ redistributive policies that support indebted households have direct employment consequences

From BHO: What Drove the Recession?

- Demand shocks (consistent with M&S 2014) caused the initial 2007–2010 employment decline
- But wage *stickiness* played a smaller role than aggregate data alone suggest
- The slow 2010–2014 recovery reflects negative labour supply shocks, not a prolonged demand deficiency

Connection to the Full Course Arc

A recurring theme: regional variation as identification, but requiring care about what is and is not identified.

- **L1 (HM19)**: Bartik instruments identify local wage gains from TFP $\rightarrow$ counterfactual aggregate welfare requires a model
- **L5 (HM24)**: shift-share for TFP shocks; indirect effects cross city boundaries $\equiv$ AII/ATI from Lecture 7
- **L6 (MRR)**: gravity structure for commuting flows; multiplier only aggregates with a structural model
- **L6 (Notowidigdo)**: housing supply elasticity for local incidence; doesn't identify aggregate housing market effects
- **L8 (M&S 2014 $\rightarrow$ BHO)**: today's payoff — the explicit methodological lesson about regional vs. aggregate inference

Forward: Lectures 9–10

Dynamic spatial models (Greaney et al.; Caliendo et al.; Guaitoli & Pancrazi) are the next frontier: static regional regressions cannot capture the intertemporal adjustment paths that matter for welfare.

What We Learned Today

M&S (2014)

- $\Delta HNW \rightarrow$ NT employment $\downarrow$; tradable unaffected
- IV: Saiz supply elasticity; two classification methods
- Little wage or migration adjustment
- *Limit*: scaling to aggregates requires same shock composition

BHO (2019)

- Wages more flexible than aggregate data suggest (76% vs. 50%)
- Regional WPC = Aggregate WPC: κ_w is scale-invariant
- 2007–10: demand shocks; 2010–14: labour supply shocks
- Regional ID blind to aggregate-only shocks

The Unifying Lesson

Regional identification is powerful for structural parameters. It cannot substitute for a structural model when shocks differ between regional and aggregate episodes.

Further Reading

Core papers:

- Mian, A. and Sufi, A. (2014). "What Explains the 2007–2009 Drop in Employment?" *Econometrica*, 82(6), 2197–2223.
- Beraja, M., Hurst, E., and Ospina, J. (2019). "The Aggregate Implications of Regional Business Cycles." *Econometrica*, 87(6), 1789–1833.

Background and supplementary:

- Mian, A., Rao, K., and Sufi, A. (2013). "Household Balance Sheets, Consumption, and the Economic Slump." *QJE*, 128(4), 1687–1726.
- Nakamura, E. and Steinsson, J. (2014). "Fiscal Stimulus in a Monetary Union." *AER*, 104(3), 753–792. *[Pioneered the structural approach to regional multipliers]*
- Saiz, A. (2010). "The Geographic Determinants of Housing Supply." *QJE*, 125(3), 1253–1296.
- Chodorow-Reich, G. (2018). "Geographic Cross-Sectional Fiscal Spending Multipliers." *AEJ: Economic Policy*, 11(2), 1–45.