

Macro-Spatial Economics

General Info

Professor: Roberto Pancrazi, University of Warwick

Email: R.Pancrazi@warwick.ac.uk

Course website: <https://www.robertopancrazi.com/macro-spatial-eui/>

Overview

This course examines how geography and location shape macroeconomic outcomes such as productivity, growth, and employment. It explores why large regional wage and productivity gaps persist despite limited worker mobility, and how spatial frictions lead to inefficiencies in the distribution of talent and capital. Students will develop a quantitative spatial equilibrium framework to analyze these dynamics and assess policies like place-based subsidies and housing regulations. Combining theory with cutting-edge empirical tools, the course connects local economic behavior to national and global outcomes, addressing topics from housing market cycles to dynamic spatial adjustment, intergenerational sorting, and optimal spatial policy.

Grading

The final grade consists of two components: **10% for class attendance** and **90% for an end-of-course referee report**. Specifically, I will select a set of recent job market papers (ideally from this season) related to spatial macroeconomics. Each student will be required to write a **referee report** on one of these papers, assessing its strengths and weaknesses, and suggesting possible improvements or extensions.

Course Schedule

Lecture 1: 08/04/2026 11:00-13:00

Lecture 2: 09/04/2026 14:00-16:00

Lecture 3: 10/04/2026 11:00-13:00

Lecture 4: 13/04/2026 08:45-10:45

Lecture 5: 15/04/2026 08:45-10:45

Lecture 6: 17/04/2026 11:00-13:00

Lecture 7: 20/04/2026 08:45-10:45

Lecture 8: 22/04/2026 11:00-13:00

Lecture 9: 23/04/2026 14:00-16:00

Lecture 10: 24/04/2026 11:00-13:00

Extended Preview

Why does a software engineer in San Francisco earn substantially more than an equally skilled engineer in Cleveland? Why do young Italians and Spaniards migrate to Northern Europe while Southern regions struggle with unemployment and stagnant wages? Why did London's productivity surge while the UK's Midlands declined, and what does this divergence mean for national growth? These patterns reveal a fundamental truth: where economic activity takes place is not a detail—it's central to understanding aggregate outcomes in productivity, growth, and employment.

This course explores how space and geography shape macroeconomic performance. Workers don't move freely to high-wage locations, firms cluster in productive cities despite congestion costs, and local housing markets can amplify or dampen national business cycles. We begin with a striking empirical puzzle: despite enormous wage and productivity differences across regions—both within countries and across Europe—spatial reallocation is surprisingly limited, leaving talent and capital inefficiently distributed. From this observation, we build a rigorous quantitative toolkit—the workhorse spatial equilibrium framework—that allows us to measure these frictions, trace their aggregate consequences, and evaluate policy interventions from place-based subsidies to housing regulations.

The course is organized around core macro questions viewed through a spatial lens: How do local productivity shocks propagate through the economy? Why do housing markets amplify business cycles? How should fiscal policy account for regional heterogeneity? Can temporary shocks have permanent spatial effects through homeownership and forward-looking location choice? Along the way, you'll master cutting-edge empirical methods—gravity equations, exact hat algebra, and structural estimation—that have become essential tools in applied macroeconomics. By the final lectures, we'll push to the research frontier: dynamic spatial equilibria with forward-looking agents, the intergenerational mechanics of sorting across cities, and the design of optimal spatial policy when worker types cluster endogenously—the live questions where the next generation of spatial macro research is being written.

This is macroeconomics for a world where location matters, where local policies have national consequences, and where understanding the spatial distribution of economic activity is key to understanding aggregate outcomes.

LECTURE 1: Why Space Matters for Macro

Goal:

Motivate the course with a compelling empirical puzzle

Core Papers:

- Hsieh & Moretti (2019) “Housing Constraints and Spatial Misallocation” *American Economic Journal: Macroeconomics*
- Greaney, B. (2026) “Housing Constraints and Spatial Misallocation: Comment” *American Economic Journal: Macroeconomics*

Structure:

- Start with facts: U.S. wage dispersion across cities, housing price growth
- The puzzle: Why don't workers move to high-wage cities?
- Walk through the model: spatial equilibrium, housing supply constraints
- The punchline: Spatial misallocation reduces aggregate GDP by 36%
- Discussion: This is a macro problem with spatial roots

Why these papers:

Hsieh & Moretti (2019) is the motivating paper for the entire course: spatial frictions — primarily housing supply constraints — prevent efficient worker reallocation and impose large aggregate welfare costs, with the headline estimate that removing constraints in high-productivity cities would raise aggregate GDP by 36%. Greaney (2026) provides the essential methodological counterpoint: the comment argues that this welfare

gain depends critically on model specification — in particular on assumptions about where displaced workers relocate and how agglomeration forces are aggregated — and that correcting these yields substantially smaller estimates. Assigning both papers together is pedagogically deliberate: students see not only the original result but also how model assumptions drive welfare conclusions, a lesson that recurs throughout the course. The Hsieh–Moretti–Greaney debate is the first instance of a deeper theme: general-equilibrium effects are hard to identify, easy to misstate, and central to spatial macroeconomics.

Supplementary Reading:

- Moretti (2011) “Local Labor Markets” *Handbook of Labor Economics* (for context)
-

LECTURE 2: Quantitative Spatial Economics – Facts, Models, and Simple Frameworks (AA Sections 1–3)

Goal:

Establish the empirical motivation for quantitative spatial models and develop the core building blocks: spatial equilibrium, gravity, and simple quantitative frameworks

Core Paper:

- Allen & Arkolakis (2025) “Quantitative Regional Economics” *Handbook of Regional and Urban Economics*, Volume 6, Sections 1–3

Structure:

- Empirical facts about the spatial economy: wage dispersion, trade flows, city sizes
- Two seminal models: Roback (1982) on spatial equilibrium and Krugman (1991) on agglomeration
- A simple quantitative model combining both traditions
- Equilibrium conditions and how gravity equations emerge naturally
- Extensions: alternative preferences, housing, congestion, multiple sectors

Why this paper:

Sections 1–3 of Allen & Arkolakis (2025) provide the empirical grounding and conceptual entry point for the entire quantitative toolkit. Students see why the facts demand a spatial framework before they are asked to build one. The Roback and Krugman models introduced here will be nested as special cases in Lectures 3 and 4.

Supplementary Reading:

- Roback, J. (1982) "Wages, Rents, and the Quality of Life" *Journal of Political Economy*, 90(6), 1257–1278
 - Rosen, S. (1979) "Wage-based Indexes of Urban Quality of Life" in Mieszkowski, P. and Straszheim, M. (eds.) *Current Issues in Urban Economics*, Johns Hopkins University Press
 - Krugman, P. (1991) "Increasing Returns and Economic Geography" *Journal of Political Economy*, 99(3), 483–499
-

LECTURE 3: Quantitative Spatial Economics – The General Framework (AA Sections 4–6)

Goal:

Develop the general workhorse model that nests all prior frameworks as special cases, and understand its equilibrium properties

Core Paper:

- Allen & Arkolakis (2025) “Quantitative Regional Economics” *Handbook of Regional and Urban Economics*, Volume 6, Sections 4–6

Structure:

- The general workhorse framework: unifying all previous models as special cases
- Equilibrium properties: existence, uniqueness, and comparative statics
- Taking the model to data: estimating trade frictions and structural elasticities
- Exact hat algebra: computing counterfactuals without full calibration
- Applications and lessons: infrastructure, trade shocks, policy evaluation

Why this paper:

Sections 4–6 build the general framework that subsumes Roback and Krugman as special cases. Students learn the conditions for existence and uniqueness of spatial equilibrium and develop comparative statics intuition. This provides the structural backbone that Lecture 4 will then operationalise empirically.

LECTURE 4: Quantitative Spatial Economics – Empirical Methods and Applications (AA Sections 7 to End)

Goal:

Take the quantitative spatial framework to data: estimate trade frictions and structural elasticities, compute counterfactuals using exact hat algebra, and evaluate policy applications

Core Paper:

- Allen & Arkolakis (2025) "Quantitative Regional Economics" *Handbook of Regional and Urban Economics*, Volume 6, Sections 7 to End

Structure:

- Estimating trade frictions and bilateral flows: gravity estimation in practice
- Exact hat algebra: computing counterfactuals without full structural calibration
- Identifying structural elasticities from data: trade costs, preferences, and agglomeration
- Applications: infrastructure investment, trade liberalisation, place-based policy counterfactuals
- What the toolkit cannot answer: welfare and incidence questions that require Lecture 5

Why these papers:

Sections 7 to end of Allen & Arkolakis (2025) complete the quantitative toolkit by showing how to operationalise the theory. Exact hat algebra is the key methodological contribution: it allows rigorous counterfactual analysis without fully calibrating the model, making the toolkit practical for policy evaluation.

This lecture is the technical capstone of the AA sequence before the course turns to welfare questions in Lecture 5.

LECTURE 5: Welfare and Incidence of Local Shocks

Goal:

Quantify the welfare gains from local productivity shocks and determine who bears them: workers, landowners, non-movers, and residents of other cities

Core Papers:

- Desmet & Rossi-Hansberg (2013) “Urban Accounting and Welfare” *American Economic Review*
- Hornbeck & Moretti (2024) “Estimating Who Benefits from Productivity Growth?” *Review of Economics and Statistics*

Structure:

- Review: spatial equilibrium welfare accounting from DR
- HM24: general-equilibrium incidence — wages, rents, and inequality
- The role of housing tenure: renters vs. homeowners face different welfare impacts
- Connecting back to HM19: the friction that drives misallocation also shapes who bears the costs
- Transition: from static incidence to dynamic adjustment mechanisms

Why these papers:

DR and HM24 answer the central welfare question that the AA toolkit (Lectures 2–4) makes possible to ask: when a city becomes more productive, how large are the welfare gains, and who captures them? DR provides the accounting framework; HM24 pushes the frontier by estimating indirect effects through migration and rent responses. Together they close the first arc of the course — we now know why space matters (Lecture 1), how to measure it (Lectures 2–4), and who bears the costs and gains (Lecture 5).

LECTURE 6: Migration, Commuting, and Labor Market Adjustment

Goal:

Analyze the mechanisms through which workers adjust to local labor demand shocks, and why adjustment is incomplete

Core Papers:

- Monte, Redding & Rossi-Hansberg (2018) “Commuting, Migration, and Local Employment Elasticities” *American Economic Review*
- Notowidigdo (2020) “The Incidence of Local Labor Demand Shocks” *Journal of Labor Economics*

Structure:

- MRR: commuting as the fast adjustment margin, migration as the slow margin
- Quantifying commuting and migration elasticities using the structural framework
- Notowidigdo: why poor regions retain population despite negative shocks
- The role of housing prices and government transfers in impeding outmigration

- Persistent local unemployment and aggregate inefficiency: the incomplete adjustment problem

Why these papers:

MRR builds the theory that explains why the migration shares used as data in HM24 look the way they do: commuting and migration operate at different frequencies and are complements, not substitutes. Notowidigdo provides the counterpart: even the slow adjustment margin is impeded by housing prices and transfers. Together the two papers answer a question left open by the first arc — given the incidence we identified, what are the adjustment mechanisms, and why do they fail?

LECTURE 7: Spatial Econometrics and Empirical Methods

Goal:

Develop the spatial econometric toolkit needed to work with geographically structured data: spatial dependence, weight matrices, spatial regression models, and spillover decomposition

Core Papers:

- Sarrias, M. (2020) “Notes on Spatial Econometrics”, Universidad de Talca (lecture notes)

Structure:

- Why spatial econometrics? Spatial dependence, spatial autocorrelation, and when standard OLS fails
- Spatial weight matrices: contiguity (Queen, Rook), distance-based, and row standardisation
- Testing for spatial autocorrelation: Moran’s I and its interpretation
- Taxonomy of spatial models: Spatial Lag (SLM), Spatial Durbin (SDM), and Spatial Error (SEM)
- Interpreting spillovers: direct vs. indirect marginal effects and partitioning global effects over space; hands-on R session with sf and spdep

Why these papers:

Sarrias (2020) provides a self-contained, R-oriented treatment of spatial econometrics that matches the level and tools of the course. Placing this lecture immediately after MRR and Notowidigdo (Lecture 6) ensures students can formalise the spatial spillover patterns they have already encountered — local employment multipliers, commuting externalities, amenity gradients — using the correct econometric framework. The SDM is the natural estimator for the spatial equilibrium regressions that appear throughout the remaining lectures.

Supplementary Reading:

- LeSage, J. and Pace, R.K. (2009) “Introduction to Spatial Econometrics” *CRC Press (reference text)*
 - Conley, T.G. (1999) “GMM Estimation with Cross-Sectional Dependence” *Journal of Econometrics*
 - Lansley, G. and Cheshire, J. (2016) *An Introduction to Spatial Data Analysis and Visualisation in R*, Consumer Data Research Centre (freely available online)
-

LECTURE 8: Business Cycles, Housing, and Regional Dynamics

Goal:

Understand how housing wealth transmits local shocks into aggregate fluctuations, and how regional variation identifies the underlying macro mechanisms

Core Papers:

- Mian, A. and Sufi, A. (2014) "What Explains the 2007–2009 Drop in Employment?" *Econometrica*
- Beraja, Hurst & Ospina (2019) "The Aggregate Implications of Regional Business Cycles" *Econometrica*

Structure:

- M&S (2014): housing net worth channel — cross-county identification strategy, tradable vs. non-tradable employment, Saiz IV
- Key result: a 10pp decline in housing net worth reduces non-tradable employment by 3.7pp; tradable employment unaffected (placebo)
- The limits of regional identification: back-of-the-envelope scaling from regional to aggregate misses aggregate-only shocks
- BHO: structural fix — regional WPC identifies wage-stickiness parameter ξ_w shared with aggregate model; observational equivalence problem
- BHO results: wages more flexible than aggregate data suggest (75% vs. 50% annual reset); demand shocks drove 2007–2010; labour supply shocks drove slow recovery 2010–2014

Why these papers:

M&S (2014) establishes the regional evidence: housing net worth collapses reduced local employment through the demand channel, with wages and workers failing to adjust. BHO (2019) then shows why extrapolating regional elasticities to aggregate conclusions is hazardous: regional regressions difference out aggregate-only shocks (labour supply, monetary policy) that were important for the recession and the slow recovery. Seeing M&S first is essential — students must understand the regional approach before they can appreciate what BHO are critiquing and what their structural methodology adds. The lesson is about the limits of regional identification, not just its power.

Supplementary Reading:

- Kaplan, G., Mitman, K. and Violante, G.L. (2020) "The Housing Boom and Bust: Model Meets Evidence" *Journal of Political Economy*
- Mian, A., Rao, K. and Sufi, A. (2013) "Household Balance Sheets, Consumption, and the Economic Slump" *Quarterly Journal of Economics*

LECTURE 9: Dynamic Spatial Models

Goal:

Introduce forward-looking spatial equilibrium models and understand why statics are insufficient for welfare analysis

Core Papers:

- Greaney, J., Parkhomenko, A. and Van Nieuwerburgh, S. (2025) "Dynamic Urban Economics" Working Paper
- Caliendo, Dvorkin & Parro (2019) "Trade and Labor Market Dynamics" *Econometrica*

Structure:

- Why statics are insufficient: transition dynamics, forward-looking migration, homeownership and capital gains

- Greaney/Parkhomenko/VNW: homeownership, housing as an asset, and forward-looking location choice
- Welfare comparisons require tracking the transition path, not just steady states
- CDP: large-scale empirical grounding — adjustment to the China trade shock takes decades
- Multiple equilibria and hysteresis: can temporary shocks have permanent spatial effects?

Why these papers:

This lecture introduces the dynamic framework that the remaining lectures build on. Students learn what forward-looking behavior and homeownership add to the static toolkit from Lectures 2–4, and why transition dynamics change welfare calculations. CDP provides large-scale empirical discipline by showing that adjustment to trade shocks takes decades. This lecture directly sets up Lecture 10, where Greaney et al. is examined in depth alongside the optimal sorting framework.

LECTURE 10: Sorting

Goal:

Connect the course’s quantitative framework to the frontier of spatial macro research, centering on the role of intergenerational dynamics and optimal sorting

Core Papers:

- Guaitoli, Pancrazi & Raimondo (2026) Working Paper
- Fajgelbaum & Gaubert (2020) “Optimal Spatial Policies, Geography, and Sorting” *Quarterly Journal of Economics*, 135(2), 959–1036

Structure:

- Guaitoli, Pancrazi & Raimondo: an OLG spatial model with intergenerational congestion and cross-country income data
- What the paper inherits from Greaney et al. (Lecture 9) and what it adds
- Fajgelbaum & Gaubert: optimal spatial taxation with sorting and agglomeration
- The normative question: when should policy reinforce or counteract spatial sorting?
- Open research questions: where is the frontier in spatial macroeconomics?

Why these papers:

The course ends with live frontier research. Guaitoli, Pancrazi & Raimondo lands naturally here because students already have the Greaney et al. dynamic framework in hand and can evaluate precisely what the new paper adds. Fajgelbaum & Gaubert closes the normative arc opened in Lecture 1: having traced misallocation, adjustment, and dynamics, the course ends by asking what optimal spatial policy looks like when sorting is endogenous.
