

A Spatial Paradox in U.S. Metropolitan Growth: Spatially Clustered Growth Rates, Spatially Independent Levels

Michael L. Chohaney

chohaney.michael@gmail.com

Oleg A. Smirnov

University of Toledo

Research Article

Keywords: regional economic growth, metropolitan economies, spatial dynamics, spatial autocorrelation, spatial Markov chains

Posted Date: September 30th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11017684/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

A Spatial Paradox in U.S. Metropolitan Growth: Spatially Clustered Growth Rates, Spatially Independent Levels

Abstract

Annual growth rates of real metropolitan gross domestic product (MGDP) are strongly clustered in space, yet the levels into which those rates accumulate show no detectable spatial autocorrelation. The puzzle is that each level is a running sum of past growth, so persistent spatial clustering in growth might be expected to appear in the geography of levels. Using MGDP for 376 U.S. metropolitan statistical areas from 2001 to 2023, a period spanning the 2008 financial crisis and the COVID-19 shock, we document and explain this dissociation. Growth is positively autocorrelated in every year of the panel (mean Moran's $I \approx 0.17$, $p < 0.05$ in all 22 years), while levels are not (mean $I \approx 0.00$, never significant). A spatial autoregressive accumulation model shows that two features of the growth process explain the divergence. Growth shocks are small relative to the accumulated stock, so each year adds only a thin spatially structured layer. More importantly, growth is mean-reverting, causing clusters to relocate rather than recur in the same places. Spatial dependence in levels becomes significant within the panel when we hold shock size and dependence constant but force the same spatial pattern to persist. Small shocks slow the accumulation, but cluster relocation extends the emergence of spatial dependence in levels from decades to centuries. A spatial Markov analysis supports the mechanism assumed by the model: growth states revert strongly toward the mean, mix within about two years, and transition with probabilities conditioned on neighboring states ($p < 10^{-13}$). Spatial clusters of MGDP growth are itinerant, forming, dissolving, and re-forming in shifting locations, leaving no lasting spatial imprint on levels.

Keywords: regional economic growth, metropolitan economies, spatial dynamics, spatial autocorrelation, spatial Markov chains

JEL Classification: R11 , O47 , R12 , C23

1 Introduction

GDP dynamics across metropolitan statistical areas (MSAs) reveal a puzzle. Annual growth rates of metropolitan gross domestic product (MGDP) are spatially clustered—high-growth metros tend to neighbor other high-growth metros, and likewise for low

growth—yet levels of MGD_P show no such spatial dependence. Because the spatial distribution of economic activity is the accumulation of its past growth, persistently clustered growth rates should, over time, imprint the same spatial pattern on levels of economic activity; however, this does not happen.

The spatial arrangement of growth is real but itinerant, shifting across the country over time. Clusters of MGD_P growth are irregular, varying in size and composition from year to year and rarely falling in the same place in consecutive years. This paper documents that pattern for U.S. metropolitan areas from 2001 to 2023 and explains it without assuming what determines economic growth.

First, we show by simulation that the paradox is expected when spatially clustered growth shocks accumulate at the rate actually observed across MSAs. Using a spatial autoregressive (SAR) data-generating process, we simulate annual growth patterns that carry the spatial dependence observed in the data, sum them over time, and measure the spatial dependence of the resulting levels. When the simulation is calibrated to the empirical distribution of MGD_P growth rates, the spatial structure of growth is not carried over to levels, and it appears in levels only under counterfactually larger shocks. Because real MGD_P growth rates are small relative to the stock of accumulated activity, the levels converge to the spatial pattern of growth so slowly that no such pattern appears over any feasible panel.

Second, we characterize the process that produces this behavior using a spatial Markov framework, which comprises the other half of the explanation. We show that metropolitan growth states (low, neutral, high) are strongly mean-reverting, that clusters of growth and decline dissipate within about two years, and that the probability of transitioning between growth states is conditioned on the growth states of neighboring metros. Clusters in each period are generated by spatial dependence in growth transitions; however, because that clustering is transient and mean-reverting, shifting across locations rather than recurring in place, it never accumulates into a persistent spatial pattern in MGD_P levels.

We emphasize at the outset that our analysis is descriptive and that we characterize the spatial-temporal structure of growth but do not identify the economic mechanism, such as spillover effects or common shocks among adjacent metros, that produces the neighbor co-movement (see Section 6).

Substantively, we document and resolve a regularity in how U.S. metropolitan economies grow. The geography of growth and the geography of accumulated economic activity are decoupled because the spatial structure of growth is transient and mean-reverting and never accumulates into a lasting spatial pattern of MGD_P levels. This finding sits alongside the standard view of spatial growth, in which agglomeration forces concentrate economic activity in core regions (Fujita and Thisse 2002; Krugman 1991) and regional growth is studied through spatially structured convergence and spatial growth regressions (LeSage and Fischer 2008; Redding and Rossi-Hansberg 2017). That literature concerns the concentration and convergence of economic activity levels, whereas our result concerns the spatial structure of growth rates, which is present and significant in every year yet never settles into a durable geography of economic activity. The year-to-year clustering of growth is, moreover, most pronounced during the 2008–2009 and 2020–2021 downturns, consistent with the regional co-movement

of metropolitan business cycles (Arias et al. 2016). To our knowledge, this dissociation between the spatial structure of growth and of levels has not been characterized in this form.

The finding also carries a methodological implication that is easy to misread. Because the spatial structure of growth relocates from year to year, spatial panels with a single, time-invariant weights matrix capture it poorly, and the natural response is a time-varying matrix W_t (Cheng et al. 2014; Qu et al. 2017). But a richer W_t does not resolve the paradox, because it respecifies which units are connected and how strongly, whereas our finding concerns where clustering materializes over a stable set of connections, not which links carry it. A fixed contiguity matrix already reproduces both facts, clustered annual growth and unclustered accumulated levels, so the transience is not an artifact of holding W constant. The relevant non-stationarity lives in the process that generates clustering, in its mean reversion and short mixing time, not in the weights, a dimension that time-varying- W estimators are not designed to address. Our contribution is accordingly empirical. We document a process whose spatial structure is transitory by construction, forming and dissolving within about two years, which reconciles strongly clustered growth with spatially independent levels. We return to this in Section 6.

The remainder of the paper proceeds as follows. Section 2 describes the data and the spatial weights. Section 3 documents the paradox. Section 4 presents the spatial-accumulation simulation that resolves it. Section 5 characterizes the underlying process using a spatial Markov analysis. Section 6 discusses implications and limitations, and Section 7 concludes.

2 Data and Spatial Weights

We measure metropolitan economic activity as total real gross domestic product for 376 U.S. metropolitan statistical areas (MSAs) over 2001–2023, expressed in natural logarithms. The series originates with the U.S. Bureau of Economic Analysis and is archived in the Federal Reserve Economic Data (FRED) database maintained by the Federal Reserve Bank of St. Louis.¹ Every MSA is observed in all twenty-three years, and annual growth rates are computed as the first difference of log levels, $g_{i,t} = \ln Y_{i,t} - \ln Y_{i,t-1}$, giving twenty-two annual cross-sections of growth (2002–2023).²

MSAs were chosen as the unit of analysis because their boundaries encapsulate a single demand market and labor pool, and a concentration of interconnected producers linked by transportation infrastructure, supply-chain linkages, and local economic policy and regulation. This choice of unit deliberately absorbs within-MSA agglomeration—the clustering of activity inside a metropolitan economy, which much of the spatial-growth literature studies (Fujita and Thisse 2002)—into the unit itself, so that what remains to be measured is the spatial dependence of growth between metropolitan economies.

¹Federal Reserve Economic Data, <https://fred.stlouisfed.org>; series with the prefix RGMP followed by the metropolitan core-based statistical area (CBSA) code. Total real GDP is used as the primary measure for its comprehensiveness and standard definition.

²The Bureau of Economic Analysis discontinued publication of metropolitan-area GDP with its 2024 release, so the analysis covers the full published lifespan of the series.

Constructing a meaningful spatial neighborhood structure for MSAs is not straightforward because most MSA boundaries are not contiguous. Nearer spatial units interact more than distant ones (Tobler 1970), and contiguity-based neighborhoods capture this more faithfully than distance-based alternatives (Smirnov 2016). Therefore, we created Thiessen (Voronoi) polygons around the geometric centroids of the boundaries of MSAs in the contiguous United States, yielding a contiguous tessellation on which spatial relations can be defined by the queen’s-contiguity criterion. This approach avoids the problems of the common alternatives, distance thresholds or k nearest-neighbor relations, which tend to over-relate the densely situated East Coast MSAs and under-relate the dispersed MSAs of the western states. The underlying contiguity graph is symmetric, has no islands, and averages 5.4 neighbors per MSA (minimum 1, maximum 10); W is its row-standardized form.

The tessellation and weights are built on metropolitan boundaries held fixed using the 2018 U.S. Office of Management and Budget (OMB) delineation. Aligning geography at a single vintage is standard for long metropolitan panels, since boundary revisions would otherwise break comparability across years. The 376 MSAs we analyze are those for which the Bureau of Economic Analysis reports a continuous real-GDP series over 2001–2023. The remaining metros in the 2018 delineation lack a complete series and are dropped to keep the panel balanced.

3 The Paradox

We measure spatial dependence with Moran’s I , computed separately for each year in (i) the cross-section of growth rates and (ii) the cross-section of log levels. Table 1 reports both, with inference based on 999 random permutations.

Growth rates are positively and significantly spatially autocorrelated in all twenty-two years of the panel, with a mean Moran’s I of 0.169 and $p < 0.05$. The spatial clustering of growth is strongest during the 2008–2009 financial crisis (Moran’s I of 0.262 and 0.218) and remained pronounced through the COVID-19 shock of 2020–2021 (Moran’s I of 0.171 and 0.162), consistent with the regional co-movement and propagation of metropolitan business cycles during downturns (Arias et al. 2016; Shibaev 2016). In contrast, MGDG levels have a mean Moran’s I of essentially zero and are not statistically significant in any year of the panel. This is the paradox (Figure 1): the spatial structure of the growth rates is absent from the levels into which those rates accumulate.

3.1 Itinerant clusters

The spatial clustering of growth shifts in location from year to year. Local indicators of spatial association (LISA) calculated for each year’s growth rates identify a handful of significant high–high and low–low clusters, but these clusters rarely appear in the same locations in consecutive years. Only about 20% of the metros in a significant high-growth cluster are still in a cluster the next year, and for low growth clusters the share is about 7%. Figure 2 tracks the high-growth core as it migrates from the upper Midwest in the mid-2000s to Texas and the northern Plains in the early 2010s, and back toward the Midwest by 2016. Because the clusters relocate quickly, the spatial

Table 1 Spatial autocorrelation of real MGDG growth rates and levels, by year. Moran’s I with permutation p -values (999 permutations)

Year	Growth rate		Level	
	Moran’s I	p	Moran’s I	p
2002	0.157	0.001	0.008	0.364
2003	0.135	0.001	0.007	0.363
2004	0.191	0.001	0.006	0.391
2005	0.186	0.001	0.007	0.396
2006	0.154	0.001	0.007	0.398
2007	0.117	0.001	0.007	0.375
2008	0.262	0.001	0.005	0.402
2009	0.218	0.001	0.003	0.427
2010	0.141	0.001	0.001	0.418
2011	0.075	0.013	-0.001	0.460
2012	0.125	0.001	-0.001	0.459
2013	0.142	0.001	0.000	0.469
2014	0.110	0.001	-0.002	0.461
2015	0.155	0.001	-0.002	0.469
2016	0.212	0.001	-0.003	0.515
2017	0.104	0.002	-0.006	0.531
2018	0.255	0.001	-0.006	0.536
2019	0.184	0.001	-0.004	0.514
2020	0.171	0.001	-0.004	0.501
2021	0.162	0.001	-0.005	0.505
2022	0.246	0.001	-0.004	0.521
2023	0.213	0.001	-0.004	0.513
<i>Mean</i>	<i>0.169</i>		<i>0.000</i>	

structure in each year’s growth never has time to carry over into the levels, the spatial counterpart of the rapid mixing established in Section 5.

3.2 The spatial scale of co-movement

The clustering documented above is strongest among first-order neighbors and decays beyond them, whether we measure spatial relations by adjacency or distance. Our primary measure is the spatial correlogram of [Cliff and Ord \(1981\)](#), which extends the contiguity structure of Section 2 from first to sixth order. For each order m , we row-standardize the weight matrix $W^{(m)}$ and compute the mean annual Moran’s I of MGDG growth. Co-movement is heavily concentrated among immediate neighbors. The order-one mean Moran’s I is 0.169 (the value in Table 1), roughly halves to 0.088 at order two, then declines to 0.073, 0.051, and 0.024 at orders three through five, reaching 0.015 by order six, about a tenth of its order-one value. Growth clustering is statistically significant in all 22 years at order one but only about half of the years by the fifth order.

The same decay appears when we bin pairs directly by the great-circle distance between the MSA centroids (Figure 3, panel b). MGDG growth co-movement is strongest within 100 km (mean Moran’s $I \approx 0.23$, significant in 20 of 22 years), declines




Fig. 1 Moran's I of real MGDP growth rates (upper series) and levels (lower series), 2002–2023

monotonically (0.14 at 100–200 km, 0.10 at 200–350 km, 0.05 at 350–550 km, 0.03 at 550–900 km), and is indistinguishable from zero beyond roughly 900 km (mean -0.02 , never significant). By both conceptions of spatial relations, then, the dependence of growth is local and decays with separation, so the finding does not depend on how neighbors are defined.

4 Resolving the Paradox: Spatial Accumulation

4.1 A spatial autoregressive accumulation model

We address the paradox with a data-generating model that makes the accumulation of spatial growth explicit. Suppose that growth in each period is governed by a spatial autoregressive process, so that the vector of growth rates $\mathbf{g}_t$ is generated as

$$\mathbf{g}_t = (\mathbf{I} - \rho\mathbf{W})^{-1} \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim N(\boldsymbol{\mu}, \sigma^2\mathbf{I}), \quad (1)$$

where ρ is the spatial dependence parameter and $(\mathbf{I} - \rho\mathbf{W})^{-1}$ is the spatial multiplier that propagates an i.i.d. shock $\boldsymbol{\varepsilon}_t$ across the neighbor network. The level of (log) activity accumulates these growth increments,

$$\mathbf{y}_t = \mathbf{y}_{t-1} + \mathbf{g}_t = \mathbf{y}_0 + \sum_{s=1}^t \mathbf{g}_s, \quad (2)$$




Fig. 2 Significant local clusters (LISA, $p \leq 0.05$) of MSA growth rates in selected years. High-high (red) and low-low (blue) clusters shift location from year to year, illustrating the transitory character of spatial growth

starting from the observed 2001 level $\mathbf{y}_0$. By construction each period's growth field carries spatial dependence ρ . The question is whether the accumulated levels $\mathbf{y}_t$ inherit it.

The mean μ affects none of the Moran's I values below because a row-standardized W , $(\mathbf{I} - \rho W)^{-1} \mu \mathbf{1} = \mu(1 - \rho)^{-1} \mathbf{1}$ is a constant vector, removed when Moran's I centers the field. Only the shock scale σ and the initial variance V_0 govern the spatial dynamics; μ enters the scenarios (§4.4) to match the observed average growth rate.

4.2 What accumulation implies analytically

Two quantities govern the levels' dependence: the limit I_∞ it approaches and the rate at which it gets there. The limit follows from the variance-weighted average below, and the rate depends on how large each period's spatially structured increment is relative to the stock already accumulated. The level is the sum of two components: the fixed initial stock $\mathbf{y}_0$ (cross-sectional variance V_0 , Moran's $I = I_0 \approx 0.008$) and the accumulated structured signal (variance ks , Moran's $I \rightarrow I_\infty \approx 0.148$). Their cross term vanishes in expectation, since the accumulated shocks are mean-zero and independent of $\mathbf{y}_0$, so the Moran's I of the level after k periods is, to a close approximation, the variance-weighted average of the two components' dependence,

$$I(\mathbf{y}_{t+k}) \approx \frac{k s I_\infty + V_0 I_0}{k s + V_0}, \quad s = \sigma^2 \bar{c}, \quad \bar{c} = \frac{1}{N} \text{tr}[(\mathbf{I} - \rho W)^{-1}(\mathbf{I} - \rho W')^{-1}]. \quad (3)$$

(3) approximates the ratio of expected numerator and denominator, $\mathbb{E}[\mathbf{x}'W\mathbf{x}]/\mathbb{E}[\mathbf{x}'\mathbf{x}]$, rather than the expectation of the ratio Moran's I itself computes, $\mathbb{E}[\mathbf{x}'W\mathbf{x}/\mathbf{x}'\mathbf{x}]$. The two coincide asymptotically, since a ratio of correlated random variables carries a numerator-denominator covariance term that this approximation




Fig. 3 Spatial correlogram of MSA growth. Mean annual Moran's I of real MGDGP growth, 2002–2023, by (a) contiguity order in the queen-adjacency graph and (b) great-circle distance band between MSA centroids. Both constructions show co-movement concentrated among near neighbors and decaying with separation. Order-one Moran's I equals the growth clustering reported in Table 1

sets aside. The long-horizon simulation (§4) confirms this gap is negligible at the horizons studied here and validates the trajectory (3) implies.

Two timescales follow from (3). The accumulated structured variance overtakes the fixed stock, $ks = V_0$, only after $k^* = V_0/s \approx 1,200$ periods, and the level dependence reaches half its limit, $I = I_\infty/2$, after $k_{1/2} = V_0(I_\infty - 2I_0)/(sI_\infty) \approx 1,100$ periods. Because the growth shocks are so small, each period contributes only a thin spatially structured layer to a large, spatially independent stock, so convergence to the limit is correspondingly slow. The simulation below quantifies this rate and confirms (3), and with it that the accumulated levels remain indistinguishable from spatially random over any feasible panel.

4.3 Calibration and simulation design

The generating process (1) has a single free input, the spatial autoregressive parameter ρ , which we estimate directly from the data. A constant-only spatial autoregressive model fitted to each year's growth rates gives an average $\rho \approx 0.33$. A growth field drawn at this ρ has a Moran's I of about 0.15, and accumulating a fresh draw of the field each period drives the Moran's I of the resulting levels toward this same value ($I_\infty \approx 0.148$) rather than amplifying it. This limit sits marginally below the observed growth Moran's I of 0.169 because ρ and Moran's I are different summaries of spatial dependence, and ρ is estimated from the full spatial covariance rather than tuned to reproduce Moran's I . This small gap is immaterial to the conclusion and, if anything, conservative because a higher limit would accelerate convergence.

4.4 Simulation scenarios

To isolate *why* spatial dependence does or does not accumulate, we tune the simulation by the spatial dependence of growth (high vs. low ρ) and the scale of the growth shocks. For scale we contrast the *empirical* moments, the mean ($\mu \approx 0.018$) and standard deviation ($\sigma \approx 0.035$) of the observed MSA growth data, against *standardized* moments $(\mu, \sigma) = (0, 1)$, i.e. unit-variance shocks far larger than the real economy experiences. Crossing dependence with scale gives four scenarios:

- *Scenario A* — high ρ ; standardized moments $(0, 1)$;
- *Scenario B* — high ρ ; empirical moments $(0.018, 0.035)$ (the observed economy);
- *Scenario C* — low ρ ; standardized moments $(0, 1)$;
- *Scenario D* — low ρ ; empirical moments.

Scenarios B and D use the growth rates observed in the real metropolitan economy, and A and C inflate the growth rates to unit variance to test whether the same growth dependence would survive accumulation under larger shocks.

These four scenarios draw a fresh spatially dependent field every period, so its clustering relocates from year to year, exactly as the observed clusters do (§3.1). That built-in relocation is held fixed across A–D and never varied. To separate the effect of relocation from the effect of scale, we add a fifth scenario that changes only persistence:

- *Scenario E* — high ρ ; empirical moments $(0.018, 0.035)$; the same spatially dependent field reused in every period (a persistent, non-relocating pattern).

Scenario E matches Scenario B in spatial dependence and shock scale but differs in that its clustering remains fixed rather than relocating. The contrast between E and B therefore isolates the role of transience with ρ and scale held fixed, a comparison we develop in §4.5. For each scenario, we simulate twenty-two periods of accumulation from the observed 2001 levels and measure the Moran’s I of the accumulated level at each step, and average over 100 replications.

4.5 Result: convergence is extremely slow with the observed growth rates

Table 2 reports the Moran’s I of accumulated levels after twenty-two periods, the length of the observed panel. Scenarios A versus B vary the scale of growth shocks while sharing high spatial dependence, and Scenarios E versus B hold dependence and scale fixed and differ only in whether the clustered pattern relocates or persists. The comparison by scale reproduces the paradox, as expected. Under counterfactually large, standardized shocks (Scenario A) each period adds a thick spatially structured layer and the accumulated levels are substantially clustered within the panel (Moran’s $I \approx 0.14$), whereas at the observed scale (Scenario B) each period’s increment is small relative to the accumulated stock, so the level contains little spatial dependence (Moran’s $I \approx 0.01$).

The comparison by transience is sharper. Holding Scenario B’s dependence and scale fixed but forcing its clustered pattern to stay in place rather than relocate (Scenario E) lifts the level’s Moran’s I to ≈ 0.05 , which is statistically significant by the

end of the panel. Scale alone does not deliver the clean null, since even at the observed shock scale a persistent pattern would carry over into levels within the span of the data. It is the year-to-year relocation of clusters that keeps the observed levels spatially independent. The low-dependence scenarios (C, D) are controls, confirming that with little clustering in growth to begin with, levels stay spatially independent under either shock scale.

Table 2 Moran’s I of accumulated levels after twenty-two periods of the spatial accumulation simulation, by scenario, over 100 replications. “Empirical” and “standardized” shock moments are defined in the text. Scenarios A–D draw a fresh innovation each period (a relocating pattern); Scenario E reuses the *same* field every period (a persistent pattern) at Scenario B’s dependence and scale, isolating transience.

Scenario	Shock moments (μ, σ)	Growth input ρ	Level Moran’s I
A	standardized (0, 1)	0.332	≈ 0.14
B	empirical (0.018, 0.035)	0.332	≈ 0.01
C	standardized (0, 1)	0.05	≈ 0.01
D	empirical (0.018, 0.035)	0.05	≈ 0.01
E	empirical (0.018, 0.035), persistent	0.332	≈ 0.05
<i>Observed data</i>		—	0.00

The scale of shocks and the transience of clusters keep levels spatially independent. The theory guarantees that spatially dependent growth eventually structures the levels it accumulates (§4), but at the observed scale this happens far too slowly to show up in a feasible panel. Real growth shocks are small relative to the spread of accumulated levels, so each year adds only a thin structured layer and the level’s dependence builds slowly. Clusters relocate from year to year (§5), so successive increments do not reinforce the same locations. Running the observed-scale path (Scenario B) forward for many centuries confirms that the levels eventually acquire the dependence the theory predicts, but only long after any period we could observe. Their Moran’s I stays near 0.01 across the twenty-two-year panel, rises to about 0.02 after a century, and reaches only about 0.03 after 250 periods, still far short of the ≈ 0.15 limit. Over any observable panel, the metropolitan economy sits in the Scenario B regime, where clustered growth produces spatially independent levels.

These two forces are separable within the accumulation model, and separating them shows that transience is not a restatement of small scale. Scenarios A–D all draw a new innovation in each period, so their spatial pattern relocates by construction, but none of them holds the location of spatial clustering constant. Scenario E shares Scenario B’s spatial dependence and its empirical shock scale but holds the spatial structure fixed every period (a persistent, non-relocating pattern), so the contrast between E and B isolates transience. When the pattern relocates, each period’s field is independent of the last, so the structured variance grows in proportion to k . When it remains constant, the same field is added every period, so the variance grows in proportion to k^2 . Substituting k^2s for ks in (3) leaves the limit unchanged but shortens every




Fig. 4 Simulated Moran's I of accumulated levels against the number of accumulation periods (one period per year), for the scenarios of Table 2. Only high spatial dependence combined with counterfactually large shocks (Scenario A) drives level clustering toward its limit within the window. At the empirical shock scale the relocating economy (Scenario B) stays near zero. Scenario E holds Scenario B's spatial dependence and scale fixed but reuses the *same* spatial field every period (a persistent, non-relocating pattern). At the empirical scale this persistent pattern rises steadily within the panel, isolating the role of transience from that of scale.

intermediate timescale to its square root. At the empirical shock scale the Moran's- I half-life therefore falls from about 1,100 years to about 33 (Scenario E, Figure 4). At that scale a persistent pattern would cross the ≈ 0.05 threshold of statistical significance by year 23, within the span of the data, whereas the relocating process crosses it only after about 520 years. The relocation of clusters, not the small scale of shocks alone, is what stretches the timescale for level dependence from decades to centuries.

5 Characterizing the Process: Spatial Markov Dynamics

The simulation shows that at the observed scale of growth, spatially dependent shocks accumulate into independent levels over any feasible horizon. This section turns to the process itself and establishes, empirically, that metropolitan growth is transient and mean-reverting, the second force behind the paradox. We classify each MSA-year into

a growth state—low (bottom 25% of that year’s growth distribution), high (top 25%), or neutral (middle 50%)—and study the transitions between states.

5.1 Mean reversion and rapid mixing

Table 3 reports the aspatial one-year transition matrix. Growth states are strongly mean-reverting, and the two tails behave almost symmetrically. A high-growth metro remains high the following year only 41.7% of the time, a low-growth metro only 37.0%, with the remainder returning toward the neutral state.³ Time spent in either tail is short and of similar length, about 1.7 years for high growth and 1.6 for low. A metro leaving a tail usually returns through the neutral state, two-thirds to three-quarters of the time, rather than crossing directly to the opposite one. Any single metro therefore traverses the full range slowly, reaching the opposite extreme only after roughly five years in either direction. The composition of the cross-section recovers much faster. After a shock pushes the shares of high-, neutral-, and low-growth metros away from their long-run 25/50/25 split, that imbalance corrects within about two years, consistent with the year-to-year relocation of clusters documented above.⁴

Table 3 Aspatial one-year transition matrix for MSA growth states. Rows are the origin state and columns are the destination state

From \ To	Low	Neutral	High
Low	0.370	0.467	0.163
Neutral	0.218	0.572	0.210
High	0.195	0.389	0.417

5.2 Spatial conditioning of transitions

We estimate transition matrices conditioned on the growth context of an MSA’s neighbors. Specifically, we use the average growth rate of an MSA’s neighbors in the origin year, classified into low, neutral, and high terciles, following the spatial Markov approach (Rey 2001). We cut the focal growth distribution at quartiles to define well-populated tail states against a broad neutral middle, but bin the conditioning neighbor-average into terciles so the three conditional matrices rest on balanced samples. Table 4 reports the full set of conditional probabilities.

A transition into a high- or low-growth state becomes more likely when the neighborhood matches it and less likely when it does not. When a low-growth MSA’s

³The high persistence of the neutral state (57.2%) reflects in part its larger size (50% of the distribution versus 25% for each tail). The informative dynamics are nonetheless in the high- and low-growth tails.

⁴Formally, the long-run distribution of the states is the 25/50/25 split of the year-relative states, and the second-largest eigenvalue of the transition matrix ($\lambda_2 = 0.222$) sets how fast the state shares return to that split after a disturbance. A deviation is multiplied by roughly λ_2 each year, so after two years only $\lambda_2^2 \approx 0.05$ of it remains. The cross-section re-balances faster than any single metro crosses the full range, because it reflects every metro adjusting at once rather than one at a time.

neighbors are high-growth, its chance of transitioning to high growth rises from 0.144 to 0.214 (two-proportion $z = 3.1$, $p = 0.002$), and a high-growth MSA’s chance of remaining high rises from 0.342 to 0.477 ($z = 4.4$, $p < 0.001$). A low-growth MSA is more likely to remain low when surrounded by a low-growth neighborhood (0.393 versus 0.323; $z = -2.4$, $p = 0.016$), and neutral- and high-growth MSAs are markedly more likely to fall to low growth when their neighbors are low-growth (0.269 versus 0.174, $z = -5.8$; and 0.274 versus 0.161, $z = -4.7$; both $p < 0.001$). The three transitions into the neutral state are unaffected by neighbors ($p = 0.99$, 0.59, and 0.47). In sum, high growth concentrates among high-growth metros and contraction among contracting ones, the two directions of comparable strength, while moves into the neutral middle show no neighborhood effect at all.

Though not every individual conditional probability in Table 4 is statistically significant, a likelihood ratio test of the homogeneity of the conditional transition matrices (Bickenbach and Bode 2003) decisively rejects the null that transitions are independent of neighbors’ states ($Q = 116.9$, $df = 12$, $p < 10^{-18}$), with the Low-to-High cell darkening as neighbor growth rises (Figure 5). We stress that this establishes spatial *conditioning*, meaning that neighbors’ growth states carry information about an MSA’s own transition, rather than a causal channel. Section 6 returns to what can and cannot be inferred about this mechanism.

Table 4 One-year transition probabilities by neighbors’ growth context, with two-proportion tests of the difference between the low- and high-growth neighbor contexts. Rows are the focal MSA’s own transition (own origin to own destination state); the “low”, “neutral”, and “high” columns are the terciles of the row standardized spatial lag, the average growth of all of an MSA’s neighbors that year, not the state of any single neighbor. Row blocks sum to one across destinations

From	To	Neighbors’ average growth			p (low vs. high)
		Low	Neutral	High	
Low	Low	0.393	0.363	0.323	0.016
	Neutral	0.463	0.475	0.464	0.985
	High	0.144	0.162	0.214	0.002
Neutral	Low	0.269	0.209	0.174	< 0.001
	Neutral	0.564	0.596	0.554	0.591
	High	0.167	0.195	0.272	< 0.001
High	Low	0.274	0.204	0.161	< 0.001
	Neutral	0.384	0.437	0.362	0.472
	High	0.342	0.359	0.477	< 0.001




Fig. 5 Growth-state transitions conditioned on neighbor context. Each panel is a one-year transition matrix (rows are an MSA’s own origin growth state, columns its destination state) for MSAs whose neighbors’ average growth, the row-standardized spatial lag, fell in the low, neutral, or high tercile that year. The panels differ, with the low-to-high cell darkening as neighbor growth rises. This is the spatial conditioning, and the homogeneity test confirms that it is statistically significant.

5.3 Robustness

This section verifies that both findings, the spatial conditioning of transitions and the growth–level dissociation, survive changes in state definition, spatial weights, neighbor count, and the exclusion of crisis years. Because the primary state classification is year-relative, it could confound genuine spatial spillover with spatially correlated common shocks. We therefore re-estimate the spatial Markov model using fixed (pooled) breakpoints, so that a “high-growth” state denotes the same absolute growth threshold in every year. The spatial conditioning survives and the homogeneity test again rejects ($Q = 90.3$, $p < 10^{-13}$). The mean growth rates associated with the three states (-2.8% , $+1.9\%$, $+6.3\%$ for low, neutral, high) confirm that the states capture economically meaningful differences.

The spatial conditioning is also robust to the number of state categories and to the construction of the spatial weights. Increasing the number of states from three (quantile-based) to five (quintile-based) leaves the homogeneity test decisive ($Q = 205.6$, $p \approx 5 \times 10^{-13}$). Replacing the Thiessen-polygon contiguity weights with a k -nearest-neighbor specification ($k = 6$), constructed from distances between MSA centroids rather than from the tessellation, leaves the result stronger ($Q = 128.5$, $p < 10^{-20}$ for three states; $Q = 211.0$, $p \approx 10^{-13}$ for five). It is also insensitive to how the neighbor context is binned. The main analysis (Table 4) sorts the neighbor-average growth into terciles so the three conditional matrices rest on balanced samples. Binning it instead at the same 25/50/25 breaks used for the focal state, so the two match, leaves the test decisive ($Q = 110.2$, $df = 12$, $p \approx 5 \times 10^{-18}$ under queen contiguity, and $Q = 126.7$, $p \approx 3 \times 10^{-21}$ under k -nearest-neighbor weights). The neighbor effect is maintained whether the context is cut by terciles or quartiles.

The paradox is also robust to the weight definition and neighbor count. Recomputed under k -nearest-neighbor weights for the eight values $k = 3, 4, \dots, 10$, growth-rate Moran’s I remains positive (mean ≈ 0.13 – 0.19) and significant in at least twenty-one of the twenty-two years at every k (with a single marginal year at $k = 5$ and $k = 6$), while the Moran’s I of MGDG levels remains near zero (mean within ± 0.03 of zero) and is significant in no year at any k . This holds across all eight values of k , confirming the dissociation is not specific to the queen tessellation.

Finally, because the panel spans two major macroeconomic shocks when the spatial clustering of growth was pronounced, the 2008–2009 financial crisis and the 2020–2021 COVID-19 disruption, (Table 1), we verify that the spatial conditioning is not a crisis phenomenon. The homogeneity test rejects in the pre-2008 sub-period ($p \approx 4 \times 10^{-7}$), in the inter-crisis years 2009–2019 ($p \approx 2 \times 10^{-7}$), in the COVID-era sub-period 2020–2023 ($p \approx 8 \times 10^{-4}$), and, most tellingly, when all four crisis years (2008–2009 and 2020–2021) are excluded altogether ($p \approx 5 \times 10^{-13}$). Therefore, the spatial conditioning of growth transitions is a persistent structural feature of the panel rather than a consequence of any single recession or the pandemic.

The results are robust to distance decay and the two crises that occurred during the panel. Excluding the crisis years (2008–2009 and 2020–2021), the mean Moran’s I of growth rates is essentially unchanged (0.161 versus 0.169 over the full panel) and the Moran’s I of levels remains 0.001, so the dissociation holds in every non-crisis year. The distance correlogram of Section 3.2 is similarly stable (each band within about 0.01 of its full-panel value). The downturns amplify the spatial clustering of growth, consistent with the peaks in Table 1, but they are not what produces the paradox, which is present in non-crisis years as well.

As a continuous complement to the discrete transitions, we look past the neighbor average. The average can hide a neighborhood’s composition, since two MSAs with the same average neighbor growth may sit in very different surroundings. To account for this, we decompose the row-standardized spatial lag into the growth contributions of an MSA’s low-, neutral-, and high-growth neighbors, defined by the same year-relative quartiles. We regress the growth of an MSA on the three contributions with MSA and year fixed effects and standard errors clustered by MSA. Each contribution is that category’s summed neighbor growth as a share of the MSA’s total neighbors, so a category absent from an MSA’s neighborhood in a given year contributes zero and the panel stays balanced. Estimated on same-year contributions, the three coefficients are positive and ordered by category (high 0.0087, neutral 0.0062, low 0.0056 per standard deviation; all $p < 0.001$; within- $R^2 \approx 0.08$). Estimated on the prior year’s contributions they weaken to 0.0041, 0.0035, and a low-neighbor coefficient indistinguishable from zero. Because row-standardized weights treat every neighbor equally regardless of economic size, we also re-estimate with each neighbor weighted by its lagged share of neighborhood MGDG, and the coefficients and their ordering are essentially unchanged (high 0.0083, neutral 0.0063, low 0.0050 same-year), indicating that the spatial conditioning reflects the growth of neighbors rather than the economic mass of a few large ones. Neighbor growth and an MSA’s own growth are determined within the same year, so we read these coefficients as spatial co-movement within each

year after common national shocks are absorbed by the year fixed effects, rather than as directional spillovers.

6 Discussion

The two analyses combine into one account of the paradox. Neighboring metros grow together, so each annual cross-section of MGDG growth is spatially clustered. But growth states revert to the mean and the cross-section re-balances within about two years, so those clusters dissolve and re-form elsewhere from one year to the next. These two actions keep the clustering out of the levels. The clusters move, and each year's growth is small next to the stock of activity already accumulated, producing a pattern that does not remain constant and cannot build up.

Transience and mean reversion each prevent the spatial growth pattern from accumulating, but in different ways. Transience stops each year's structured increment from landing in the same place twice, so the spatial pattern scatters instead of compounding. Mean reversion does something stronger, because a metro's turn at either extreme is short-lived. High-growth spells last about 1.7 years and low-growth spells about 1.6 before the metro reverts, most often toward the middle and less often to the opposite tail (Table 3). A metro's earlier contribution to a cluster is not just left behind as the cluster migrates. It is overwritten by that same metro's later ordinary or below-average years (§5).

Two limitations concern how we measured the neighborhood. Throughout, the conditioning runs through the average growth of neighbors and the row-standardized spatial lag but a neighborhood's composition carries information the average discards. The decomposition in Section 5 shows that a metro's own growth rises with the growth of high-, neutral-, and low-growth neighbors alike, most steeply with high-growth neighbors, and all three associations fade sharply a year later. The fact that this effect is ordered by neighbor type and is mostly gone within a year aligns with the shifting, mean-reverting process we document. It is also why we read the same-year estimates as within-year co-movement rather than directional spillover, since a neighbor's growth and an MSA's own growth are jointly determined within the year.

Information was also lost in the spatial Markov analysis. Sorting continuous growth into three or five states throws away within-state variation, so that a metro growing at four percent and one at twelve are both just "high." The continuous fixed-effects decomposition avoids that loss, and the conditioning holds under both discretizations, but the transition matrices themselves discard detail a fully continuous treatment would keep.

We have not exhausted what the composition allows. Threshold or contagion effects, frontier effects along the borders between expanding and contracting regions, and any real asymmetry between upward and downward association require its own design. Each also bears on the spillover-versus-common-shock question, since genuine spillover should scale with the number of transmitting neighbors while a common regional shock acts through their average. Identifying the economic mechanism behind the co-movement, whether distance-weighted influence that decays with proximity or

weights built from economic linkages rather than geography, is the natural next step from documenting to explaining these dynamics.

These refinements all concern the form of the neighbor relationship at a fixed point in time, and whether that structure itself changes over time is a separate question where a moving cluster pattern is easily misread. Because the spatial organization of metropolitan growth visibly relocates from year to year (Figure 2), a natural response is to reach for a model whose structure is allowed to evolve, such as dynamic or endogenous weight matrices (Qu et al. 2017; Kuersteiner and Prucha 2020; Cen et al. 2026) or regime switching and smooth-transition specifications (Lim 2016; Mattsson and Lyhagen 2025). In our setting that response would mislead, because those methods are designed for dependence whose structure genuinely evolves, through links that appear or vanish or a dependence parameter that drifts. Here the growth field relocates under a fixed weight matrix and a constant dependence parameter, its pattern set by where each year’s shocks land rather than by any drift in the structure. An evolving-structure model fitted to it would therefore mistake ordinary year-to-year sampling variation for structural change, when the relocation is exactly what a fixed, stationary process produces.

This does not rule out the presence of spatial spillovers. The observed clustering is consistent with two mechanisms that our design cannot distinguish, reflecting the long-standing distinction between substantive and nuisance dependence (Anselin 1988; Anselin and Rey 1991). Substantive dependence refers to genuine spatial spillovers, in which growth in one metro affects growth in its neighbors. Nuisance dependence arises when neighboring metros are exposed to the same underlying shocks, such as those associated with a common industrial base. Covariates may provide a more effective way to distinguish these mechanisms than any choice of spatial weights matrix. If exposure to common shocks can be measured directly, for example through shift-share exposure to national industry shocks, then any spatial dependence that remains after accounting for that exposure becomes a more credible candidate for a genuine spatial spillover effect.

7 Conclusion

M GDP growth rates are spatially clustered in every year of the 2001–2023 panel, yet the M GDP levels into which those growth rates accumulate show no detectable spatial autocorrelation. We have shown that this apparent contradiction is the expected outcome of accumulating relatively small, spatially dependent growth shocks that are both transient and mean-reverting. A spatial autoregressive accumulation model reproduces the paradox at the empirically observed scale of growth, while a spatial Markov analysis confirms the properties assumed by the model. Growth states are mean-reverting, mix within about two years, and transition with probabilities that are significantly conditioned by neighboring growth. Because a metro’s high- and low-growth spells are short and tend to revert toward the middle, and because the resulting clusters shift from year to year, the spatial structure of growth does not persist long enough to accumulate into a corresponding pattern in M GDP levels.

Although the present application concerns metropolitan gross domestic product, the underlying point is not specific to regional economics. When a level is analyzed alongside its first differences, the level is simply the cumulative sum of those differences. This is an arithmetic identity, but it does not imply that spatial dependence in the differences will carry over to the levels. The spatial structure of changes can therefore differ substantially from that of the levels to which they accumulate, even over a finite period. The same issue arises in other stock-flow relationships, including population and net migration, employment and job flows, and capital and investment. Spatially dependent flows can accumulate into stocks with little or no spatial dependence. More generally, when working with spatial panel data, it is worth asking whether spatial dependence is present in the levels, in the changes, or only temporarily in the changes. These cases call for different approaches to modeling spatial dependence.

The deeper implication is that recognizing dynamic spatial dependence is not enough on its own, because the behavior of the clusters themselves warrants investigation. The transience that keeps strongly clustered growth from imprinting its spatial pattern on the stock it accumulates was observed during an exploratory analysis of the spatial dynamics. Assessing where growth clusters appear, how long they persist, and where they travel next is, we suggest, as important to understanding regional growth as fitting a spatial model. The exploratory, dynamics-first posture taken here is a natural complement to confirmatory spatial modeling when the shifting behavior of a spatial process is itself the object of study.

References

- Anselin L (1988) *Spatial Econometrics: Methods and Models*. Kluwer Academic Publishers, Dordrecht
- Anselin L, Rey S (1991) Properties of tests for spatial dependence in linear regression models. *Geographical Analysis* 23(4):112–131. <https://doi.org/10.1111/j.1538-4632.1991.tb00228.x>
- Arias M, Gascon C, Rapach D (2016) Metro business cycles. *Journal of Urban Economics* 94:90–108. <https://doi.org/10.1016/j.jue.2016.05.005>
- Bickenbach F, Bode E (2003) Evaluating the Markov property in studies of economic convergence. *International Regional Science Review* 26(3):363–392. <https://doi.org/10.1177/0160017603253789>
- Cen Z, Chen Y, Lam C (2026) Inference on dynamic spatial autoregressive models with change point detection. *Journal of Business & Economic Statistics* <https://doi.org/10.1080/07350015.2025.2572768>, advance online publication
- Cheng T, et al (2014) A dynamic spatial weight matrix and localized space–time autoregressive integrated moving average for network modeling. *Geographical Analysis* 46(1):75–97. <https://doi.org/10.1111/gean.12026>

- Cliff A, Ord J (1981) *Spatial Processes: Models and Applications*. Pion, London
- Fujita M, Thisse JF (2002) *Economics of Agglomeration: Cities, Industrial Location and Regional Growth*. Cambridge University Press, Cambridge, UK
- Krugman P (1991) Increasing returns and economic geography. *Journal of Political Economy* 99(3):483–499. <https://doi.org/10.1086/261763>
- Kuersteiner G, Prucha I (2020) Dynamic spatial panel models: Networks, common shocks, and sequential exogeneity. *Econometrica* 88(5):2109–2146. <https://doi.org/10.3982/ECTA13660>
- LeSage J, Fischer M (2008) Spatial growth regressions: Model specification, estimation and interpretation. *Spatial Economic Analysis* 3(3):275–304. <https://doi.org/10.1080/17421770802353758>
- Lim U (2016) Regional income club convergence in US BEA economic areas: A spatial switching regression approach. *The Annals of Regional Science* 56(1):273–294. <https://doi.org/10.1007/s00168-015-0739-0>
- Mattsson I, Lyhagen J (2025) Modeling spatial regimes with smooth transitions. *International Regional Science Review* 48(1):38–61. <https://doi.org/10.1177/01600176241237180>
- Qu X, Lee LF, Yu J (2017) QML estimation of spatial dynamic panel data models with time varying spatial weights matrices. *Journal of Econometrics* 197(2):173–201. <https://doi.org/10.1016/j.jeconom.2016.11.004>
- Redding S, Rossi-Hansberg E (2017) Quantitative spatial economics. *Annual Review of Economics* 9:21–58. <https://doi.org/10.1146/annurev-economics-063016-103713>
- Rey S (2001) Spatial empirics for economic growth and convergence. *Geographical Analysis* 33(3):195–214. <https://doi.org/10.1111/j.1538-4632.2001.tb00444.x>
- Shibaev SS (2016) Recession propagation in small regional economies: Spatial spillovers and endogenous clustering. Queen’s Economics Department Working Paper 1369, Queen’s University, Department of Economics, Kingston, Ontario
- Smirnov O (2016) Geographic space: An ancient story retold. *Transactions of the Institute of British Geographers* 41(4):585–596. <https://doi.org/10.1111/tran.12141>
- Tobler W (1970) A computer movie simulating urban growth in the Detroit region. *Economic Geography* 46:234–240. <https://doi.org/10.2307/143141>