CITIES AS PORTFOLIOS*

Ailin Zhang†

March 13, 2026

Abstract

This study treats cities as portfolios of industries and applies a portfolio framework to jointly evaluate wage growth and time-series wage risk. Using U.S. data from 1990 to 2020, it maps cities into a growth–risk space and constructs an empirical efficient frontier, showing that larger cities lie closer to this frontier primarily due to superior risk efficiency. A novel symmetric decomposition reveals that large cities' growth advantage mainly reflects more favourable initial industry composition, while their risk advantage is driven predominantly by lower city-specific within-industry variance rather than cross-industry covariance, with reallocation playing a secondary role. To shed light on the mechanisms suggested by the decomposition, the analysis exploits U.S. trade liberalization with China as a spatially varying exogenous shock. The results show that cities with greater skill portability and more quality firms experience muted wage declines. Taken together, the findings indicate that urban growth and risk outcomes reflect the interaction between inherited portfolio structure and local adjustment capacity, with the latter governing how industry shocks propagate over time.

Keywords: urban growth and risk; wage volatility; industry composition; shock transmission.

*I am indebted to my advisors Gabriel Ahlfeldt and Christian Hilber for their guidance and support. I thank Sara Bagagli, Xavier Jaravel, Henry Overman, Giorgio Pietrabissa and seminar participants at LSE for discussions.

†London School of Economics and Political Science. Email: a.zhang12@lse.ac.uk.

1 INTRODUCTION

Cities differ in how efficiently they balance wage growth and risk. Some sustain rapid wage growth with little volatility, while others experience large swings with limited long-run gains. These differences become particularly salient during major aggregate shocks. When the 2008 financial crisis hit, wages in Austin and Denver dipped mildly and volatility soon stabilized, whereas in Detroit and Cleveland they fell sharply and remained unstable long after the recovery. What determines these stark differences in cities' risk-return performance?

Classic urban theories such as Marshallian agglomeration and labour market pooling highlight how large cities achieve higher productivity and pool idiosyncratic firm-level risk through thicker markets. Yet this emphasis on "urban advantages" leaves two facts insufficiently explained. First, stability varies widely even among similarly large cities, with some exhibiting persistently high wage volatility. Second, existing evidence centres on individual- or firm-level risk pooling, with limited attention to two inherently aggregate margins operating at the city–industry level: industry-wide shocks that affect all firms within a sector, and the co-movement of shocks across industries within a city. These forces generate systematic wage risk that cannot be diversified through micro-level mechanisms alone. Together, these limitations make it difficult to infer city-level resilience from standard agglomeration mechanisms alone, motivating an approach that evaluates wage growth and wage risk jointly and links aggregate volatility to cities' industry-portfolio structure.

This paper views urban labour markets as portfolios of industries, akin to portfolios of assets in financial markets. In this framework, city growth reflects both industry-level growth and the evolution of industry weights, while aggregate wage volatility depends not only on the fluctuations of individual industries but also on the covariance structure linking industries within a city. Unlike firm-level risk pooling, which primarily mitigates firm-level idiosyncratic shocks, this portfolio perspective directly captures how industry-wide shocks and their co-movement shape aggregate urban wage risk. As in financial portfolio analysis, city performance cannot be summarized by wage growth or wage volatility in isolation, but by how efficiently risk is transformed into return. We therefore evaluate cities in a joint growth–risk space, analogous to the mean–variance framework in finance, and measure efficiency by each city's distance to an empirical efficient frontier. This portfolio-based perspective provides a unified language for measuring, decomposing, and comparing cities' growth–risk profiles, and helps explain both large cities' growth–risk advantages and the persistence of substantial differences in wage volatility among cities of similar size.

The portfolio lens also helps organize several long-standing urban debates that are often treated separately. First, cross-city differences in industrial composition can be interpreted as differences in portfolio strategy—diversification versus specialization—so the classic urbanization-versus-localization question naturally becomes a return–risk question rather than a return-only one. Second, it reframes the tension between structural path dependence and local adjustment: as in finance, where realized performance reflects both passive exposure and active portfolio adjustment, cities with similar initial industry shares can nonetheless exhibit divergent growth–risk outcomes through differences in subsequent industry reallocation and shock propagation. Third, it clarifies the limits of sectoral classifications as sufficient statistics for local economies. As in finance, where securities within the same asset class differ in their risk and return profiles, an "industry" can embed distinct firm compositions across cities, generating heterogeneous sources of growth and risk.

The analogy to financial portfolios is informative but not literal, and its limits are central to interpretation. First, unlike financial investors, cities do not freely choose or trade their "assets": industry composition emerge from firms' location decisions and local equilibrium forces rather than deliberate portfolio choice. Second, adjustment is also more costly and constrained, as reallocating activity across industries requires firm entry and exit, worker mobility under institutional frictions, making rapid reallocation far more limited than in financial markets. Third, industries are not standardized assets: the same industry can embed different firm compositions and technologies across cities, and there is no centralized market in which such claims are traded. Taken together, these features imply that city performance reflects both industrial composition and locally mediated adjustment, motivating a conceptual framework that measures efficiency without interpreting it as frictionless optimization outcome. Finally, while this paper adopts an industry-based perspective, it abstracts from other dimensions of urban performance—such as amenities and housing markets—which may also shape cities' growth–risk profiles and are natural extensions for future work.

The first part of the paper maps U.S. cities (CBSAs) into a growth–risk space by plotting long-run wage growth against wage-growth volatility over 1990–2020, making the urban–finance analogy operational. In this framework, industry wage growth plays the role of asset returns, while citywide wage volatility reflects portfolio risk arising from both within-industry fluctuations and cross-industry co-movement. Motivated by the mean–variance framework of Markowitz (1952), I construct an empirical urban efficient frontier as the set of cities achieving the highest wage growth for a given level of wage risk, and measure each city's efficiency by its distance to this frontier. Two patterns emerge. First, larger cities are consistently closer to the efficient frontier in every decade. Second, the density gradient is more robust and economically meaningful for risk efficiency than for growth efficiency.

I then develop a symmetric decomposition framework—akin to performance attribution in finance (Brinson et al., 1986)—that decomposes both city growth and city risk into benchmark exposure, industry reallocation, and an idiosyncratic residual, isolating the key margins underlying large cities' efficiency advantage. The decomposition shows that large cities' growth advantage primarily reflects a more favourable initial industry mix. In contrast, their risk advantage is driven mainly by lower idiosyncratic industry risk, with reallocation playing a secondary role. Further analysis indicates that this lower idiosyncratic risk arises largely from a smaller aggregate contribution of within-industry volatility rather than from cross-industry covariance. While diversification across industries mechanically plays a central role in reducing the aggregate contribution of within-industry volatility, within-industry characteristics—such as the number of establishments and the prevalence of medium-to-large establishments—are also associated with lower risk. These decompositions are informative about *where* large cities' efficiency advantage comes from, but they do not by themselves establish *why* some cities are better able to reallocate activity and absorb shocks, nor do the cross-sectional correlations pin down causal mechanisms. Moreover, diversification and industry composition are long-run equilibrium outcomes that primarily capture cities' baseline exposure to shocks, rather than their short-run adjustment capacity, making them difficult to interpret as causal mediators in a short-run shock setting.

The second part of the paper therefore shifts from accounting to mechanisms, focusing on city characteristics that govern adjustment capacity along the two margins motivated by the decomposition: skill portability, which facilitates employment reallocation across industries, and the prevalence of medium-to-large establishments, which might buffer within-industry risk. I test these hypotheses

by exploiting U.S. trade liberalization with China, triggered by the granting of PNTR (Permanent Normal Trade Relations), as an exogenous shock to local wage dynamics, using city-level exposure measures constructed from the data in Pierce and Schott (2020). A dynamic difference-in-differences approach reveals that by 2010, an interquartile-range (IQR) increase in skill portability attenuates the wage decline associated with an IQR increase in PNTR exposure by approximately 1.8 log points, offsetting roughly 49% of the baseline aggregate wage loss. A higher pre-shock prevalence of medium-to-large establishments also buffers the shock, preserving about 1.0% of local wages and insulating workers from 27% of the baseline effect. Complementary evidence from Compustat reveals that this within-industry resilience is fundamentally rooted in the micro-level financial and operational advantages of corporate anchors, such as lower leverage and higher interest coverage. Mediation analysis further demonstrates that these two structural channels account for over 75% of the total scale-dependent resilience, effectively opening the "black box" of the urban size advantage. By contrast, other city characteristics such as number of firms and industry diversification do not exhibit statistically meaningful effects on wage responses to the China shock. This pattern suggests that initial portfolio structure (industry mix and diversification) shapes cities' long-run wage risk through baseline exposure, while adjustment capacities related to human capital and firms govern the short-run transmission of industry shocks. Observed urban outcomes reflect the interaction of these two forces over time.

The paper contributes to several strands of the literature. First, it adds to the large body of work on Marshallian agglomeration economies, labour market pooling and matching externalities (Marshall, 1890; Duranton and Puga, 2004). To our knowledge, this is the first study to adopt an urban portfolio perspective and to provide a unified framework that jointly quantifies the sources of wage growth and wage risk, treating wage volatility as a margin of economic importance comparable to mean growth. Existing studies explaining the volatility advantage of large labour markets mainly emphasize more efficient worker–firm matching and lower adjustment costs (Conte et al., 2024; Moretti and Yi, 2024), or the diversification of firms' idiosyncratic productivity shocks (Krugman, 1991; Overman and Puga, 2010; Kikuchi and O'Connor, 2025). Whereas most existing studies emphasize firm-level heterogeneity, this paper shifts attention to industrial composition as the key determinant of urban growth and risk outcomes. In this respect, it is closely related to Heblich et al. (2025), who documents the long-run productivity costs of early industrial specialization in British cities, showing that a 0.1 increase in the Herfindahl index in 1881 is associated with about a 24% reduction in productivity. While that study highlights the importance of initial industry composition and path dependence—akin to initial portfolio weights—this paper further shows that, beyond initial composition and diversification, cities' long-run performance crucially depends on adjustment margins such as skill portability and firm quality, which shape cities' responses to adverse industrial shocks and resemble active portfolio adjustment in finance.

Second, this paper contributes to the literature on the determinants of wage risk. Under incomplete markets and borrowing constraints, uninsurable wage risk plays a central role in shaping individuals' life-cycle savings, consumption, and human capital investment (Levhari and Weiss, 1974; Aiyagari, 1994; Ganong et al., 2025), and can also amplify aggregate fluctuations in response to economy-wide shocks (Debortoli and Galí, 2024). Despite its importance, the existing literature has focused predominantly on wage and earnings risk associated with individual characteristics such as age and gender (Gottschalk et al., 1994; Gottschalk and Moffitt, 2009; Guvenen et al., 2017), while paying relatively limited attention to the geographic dimension of wage risk. In particular, how heterogeneity in local labour market structure generates systematic, place-based wage volatil-

ity remains understudied. Recent urban models typically conceptualize uninsurable income risk as individual-level productivity shocks evolving over the life cycle (Greaney et al., 2025). Even in frameworks that explicitly incorporate location-specific income risk, such risk is usually taken as an exogenous input rather than as an endogenous outcome of local economic structure and adjustment processes (Ortalo-Magné and Prat, 2016; Hizmo, 2023), leaving the spatial origins of wage risk dynamics largely unexplored.

Third, this paper relates to an emerging urban-finance literature that views location choice as an intertemporal investment or self-insurance decision (Ortalo-Magné and Prat, 2016; Morten, 2019; Bilal and Rossi-Hansberg, 2021; Anstreicher, 2024). The closest related study is Bilal and Rossi-Hansberg (2021), who model locations as "assets" that entail current housing costs and future payoffs and whose positions can be adjusted through migration across places. We complement this perspective by emphasizing that locations differ not only in expected growth opportunities but also in the risk associated with those returns. Viewing cities as portfolios of industries, our framework implies that individuals choosing a location effectively hold a share of a local portfolio, exposing them to a distinct combination of expected growth and risk shaped by the city's industrial structure and its evolution over time.

The remainder of the paper is organized as follows. Section 2 maps cities into a growth–risk space by viewing them as portfolios of industries, evaluates growth–risk efficiency, and decomposes wage growth and risk using a symmetric attribution framework. Section 3 describes the data used in the descriptive analysis and the event study. Section 4 exploits the China trade shock to examine how city characteristics shape wage dynamics. Section 5 concludes.

2 MAPPING CITIES INTO THE GROWTH–RISK SPACE

2.1 Cities as portfolios

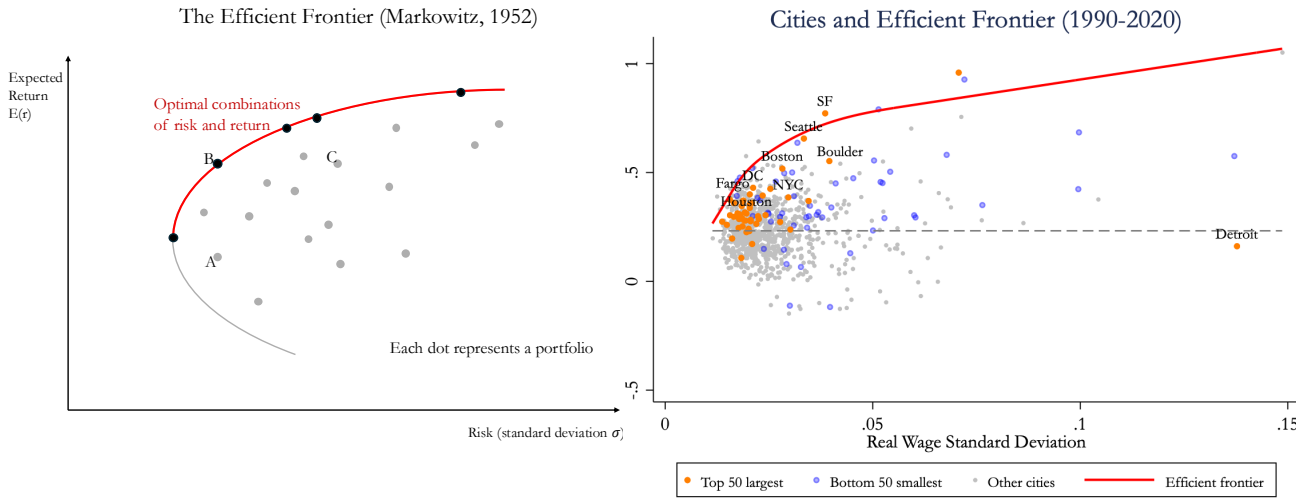


Figure 1: Efficient frontiers in finance and urban world

Figure 1 introduces the central analogy of the paper by mapping both financial portfolios and

cities into a common growth–risk space. The left panel illustrates the canonical mean–variance framework in finance (Markowitz, 1952). Each point represents a feasible portfolio constructed from the universe of risky assets, with expected return on the vertical axis and return volatility on the horizontal axis. The efficient frontier traces the upper envelope of the opportunity set: it consists of portfolios that achieve the highest expected return for a given level of risk, or equivalently, the lowest risk for a given level of return. The right panel applies the same mapping to U.S. core-based statistical areas (CBSAs) using data from Quarterly Census of Employment and Wages (QCEW), plotting cumulative real wage growth against the time-series standard deviation of real wage growth over 1990–2020. The urban efficient frontier is constructed empirically by tracing the upper envelope of the city cloud: cities are sorted by risk, record-high growth observations define an outer frontier, this is augmented with a second near-frontier layer for robustness, and a smooth cubic-spline curve is fitted through the resulting hull points, which together define the set of empirically efficient cities.

The analogy becomes natural once cities are viewed as portfolios of industries. Industries correspond to assets, and a city's employment shares correspond to portfolio weights. Aggregate wage growth reflects the weighted average of industry-level growth, while aggregate wage risk depends jointly on within-industry volatility and cross-industry covariance. As in finance, diversification reduces aggregate risk only to the extent that shocks are imperfectly correlated across industries. Cities therefore occupy distinct positions in the growth–risk space, depending on both their exposure to heterogeneous industry assets and the structure of co-movement across industries. As illustrated in Figure 1, cities such as Seattle and Boston lie close to the urban efficient frontier, combining strong long-run growth with moderate risk, whereas cities such as Detroit lie well inside the frontier, exhibiting high wage volatility without commensurate returns. Interestingly, comparing the 50 largest cities (orange) with the 50 smallest (blue) reveals a clear pattern: large cities lie systematically closer to the efficient frontier, while smaller cities are more dispersed and farther from it.

Importantly, the portfolio analogy is descriptive rather than normative. Cities do not optimally choose or freely rebalance portfolios as financial investors do; industrial composition evolves through firm dynamics, worker mobility, and local adjustment frictions. The value of the analogy lies in elevating wage volatility to a first-order object of comparison and in providing a coherent way to trace its sources. Wage volatility is economically salient because it is the risk borne by workers: by locating in a city, they pay housing and mobility costs in exchange for expected wage growth, while remaining exposed to fluctuations in that wage path—much like investors holding a risky asset. Evaluating cities in a joint growth–risk space offers a unified language for comparing urban performance and sets the stage for the decomposition of growth and risk into exposure, reallocation, and idiosyncratic components developed in subsequent sections. See Appendix A for details on the rationale and empirical implementation of the efficient frontier.

2.2 Measuring City Portfolio Efficiency

To move beyond visual comparisons and quantify the growth-risk efficiency of city portfolios, this paper introduces three related distance-based measures, drawing on insights from Kandel and Stambaugh (1995), Li et al. (2003), and Basak et al. (2002). As shown in the left panel of Fig 2, the first is an overall measure—the efficiency gap—defined as the shortest Euclidean distance from a city to the efficient frontier in the normalized growth–risk space, with smaller values indicating higher growth–risk efficiency. Two complementary directional measures characterize where this inefficiency

comes from. The growth gap is the vertical distance to the frontier and captures, for a given level of wage risk, how much growth a city forgoes relative to its most efficient counterpart. The risk gap is the horizontal distance and measures, for a given level of wage growth, the excess volatility borne relative to the most efficient benchmark. While these two gaps are mechanically correlated, they emphasize distinct dimensions of inefficiency—foregone growth versus avoidable risk. Compared with conventional portfolio efficiency measures in finance, such as the Sharpe ratio, these distance-based metrics do not require a risk-free asset, making them particularly well suited to the urban context.

A natural question is whether large cities operate more efficiently than small ones when evaluated in a joint growth–risk space. While the agglomeration literature primarily documents scale advantages in levels or growth, it remains largely silent on how growth and volatility jointly scale with city size. In the right panel of Figure 2, I address this question by estimating a simple cross-sectional regression for the period 1990–2020, relating each efficiency measure to initial city size:

$$EffiGap_c = \beta \ln Pop_c + \varepsilon_c. \tag{1}$$

I then plot the estimated coefficient β for each decade and for each of the three efficiency gaps. A negative β indicates that larger cities lie closer to the efficient frontier, while a positive β implies greater size inefficiency. For all three gap measures, the coefficient on initial log population is negative and statistically significant. Larger cities therefore tend to be closer to the growth–risk frontier. See Appendix A for more details.

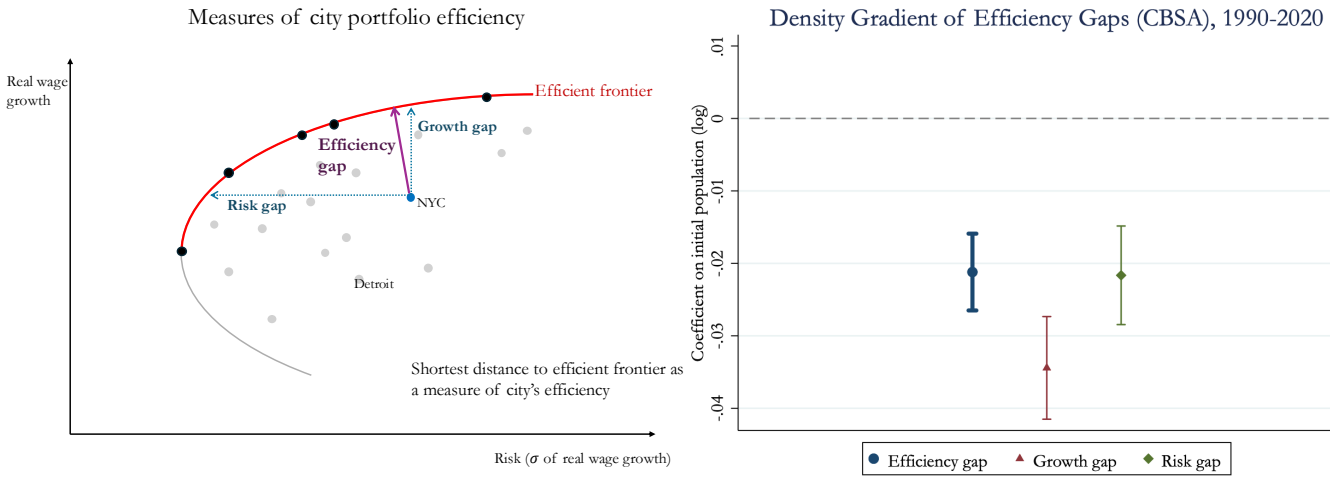


Figure 2: Efficiency measures and density gradients of efficiency gaps. The left panel illustrates three efficiency measures in normalized growth–risk space: the overall efficiency gap (Euclidean distance to the efficient frontier), the growth gap (vertical distance), and the risk gap (horizontal distance). The right panel reports coefficients from cross-sectional regressions of each efficiency measure on initial log city population.

2.3 Decomposing City Growth and Risk

2.3.1 Decomposition analysis

Equipped with our measure of portfolio efficiency, we unpack the drivers behind cities' growth and risk. A broad urban literature emphasizes that larger and more diversified cities benefit from agglomeration economies, thicker labour markets, knowledge spillovers, and a wider set of risk-sharing and adjustment margins that foster both higher growth and greater resilience. Yet these mechanisms do not readily deliver a quantitative accounting of how much of a city's subsequent performance reflects (i) its inherited industry composition—i.e., path dependence—versus (ii) the endogenous re-weighting of industries in response to shocks and evolving opportunities. A key contribution of our viewing cities as industry portfolios is that it allows us to import the portfolio performance attribution framework from finance (Brinson and Fachler, 1985; Brinson et al., 1986), providing a transparent quantitative structure to evaluate city portfolios' performance and to separate baseline composition effects from dynamic reallocation and idiosyncratic components.

Before moving to the decomposition, it is useful to note that, at any point in time t, the expected return (i.e., real wage growth) $G_{c,t}$ and variance (i.e., the variance of real wage growth) $\sigma^2_{c,t}$ of a city portfolio are given by

$$G_{c,t} = \sum_i s_{i,c,t}\, g_{i,c,t} \quad \text{and} \quad \sigma^2_{c,t} = s'_{c,t}\, \mathbf{\Omega}_{c,t}\, s_{c,t}, \tag{2}$$

where $s_{i,c,t}$ denotes the employment share of industry $i \in J$ in city c at time t, and $g_{i,c,t}$ is the expected growth of industry i at time t, which can be decomposed into a national common component and a city–industry-specific component, $g_{i,c,t} = g^N_{i,t} + \tilde{g}_{i,c,t}$. $s_{c,t}$ is the vector of industry employment shares in city c at time t, and $\mathbf{\Omega}_{c,t}$ is the $J \times J$ covariance matrix of industry real wage growth, whose diagonal elements capture industries' own variances and whose off-diagonal elements capture pairwise covariances across industries. We further allow $\mathbf{\Omega}_{c,t}$ to consist of a national common component, reflecting industry-level variance–covariance patterns at the national level, and a city-specific component that captures deviations from this national benchmark: $\mathbf{\Omega}_{c,t} = \mathbf{\Omega}^N_t + \tilde{\mathbf{\Omega}}_{c,t}$.

We now turn to the evolution of city portfolios over time and decompose real wage growth and risk using an approach borrowing the insights from the portfolio performance attribution framework pioneered by the Brinson model (Brinson et al., 1986). [1] The cumulative real wage growth over the

[1] Three remarks clarify the relationship between our decomposition and the Brinson framework. First, our approach is inspired by Brinson-style performance attribution but is not identical: the original Brinson model decomposes *returns* rather than *risk*. Risk-adjusted attribution methods exist (e.g., Ankrim, 1992), but they typically build on CAPM-style beta and link systematic risk to expected returns. In the urban setting, however, risk need not be priced in the same way because industry "assets" are not frictionlessly traded. We therefore decompose wage-growth risk symmetrically to wage growth itself. Second, whereas Brinson benchmark weights are interpreted as long-run policy weights chosen by an investment manager, our benchmark uses initial industry shares, reflecting that cities do not commit to an explicit long-run portfolio policy. Third, for parsimony we combine the Brinson *selection* component, $s_{i,c,0}(g_{i,c,t} - g^N_{i,t})$, and the *interaction* component, $(s_{i,c,t} - s_{i,c,0})(g_{i,c,t} - g^N_{i,t})$, into a single idiosyncratic-growth term. This is consistent with the view that the interaction term is not a distinct margin of active decision-making and is often grouped with selection in practice (Campisi, 2004; Bacon, 2019).

horizon T can be decomposed into:

$$G_{c,T} = \underbrace{\sum_{i,t} s_{i,c,0}\, g^N_{i,t}}_{\text{Benchmark growth}} + \underbrace{\sum_{i,t} (s_{i,c,t} - s_{i,c,0})\, g^N_{i,t}}_{\text{Reallocation gain}} + \underbrace{\sum_{i,t} s_{i,c,t}\, (g_{i,c,t} - g^N_{i,t})}_{\text{Idiosyncratic growth}}, \tag{3}$$

where the first term captures *benchmark growth*, defined as the cumulative wage growth a city would experience if it retained its initial industry employment shares at $t = 0$ while each industry grew at its national rate $g^N_{i,t}$. This component corresponds to the *passive return* in portfolio performance attribution, reflecting the contribution of initial portfolio weights evaluated at benchmark returns; for example, cities such as San Jose or San Francisco, with large initial exposure to nationally fast-growing technology and professional service sectors, exhibit high benchmark growth even absent local adjustment. The second term captures the excess return or loss generated by changes in employment shares, $s_{i,c,t} - s_{i,c,0}$, across industries $i \in J$, evaluated at national industry growth rates over the period. This component represents an active return analogous to timing or tactical asset allocation, reflecting cities' ability to reallocate labour across industries subject to skill specificity and adjustment frictions. Manufacturing-intensive cities such as Detroit, where human capital is highly industry-specific and reallocation costs are high, exhibit limited rebalancing, whereas cities like Austin might face lower adjustment frictions and more swiftly shift employment toward expanding sectors. The third term captures *idiosyncratic growth*, defined as within-industry performance relative to national benchmarks $g_{i,c,t} - g^N_{i,t}$ and weighted by contemporaneous industry shares, and is directly analogous to *security selection* in portfolio performance attribution—namely, how assets within a given sector perform relative to their benchmark index.[2] Its sign and magnitude are a priori ambiguous and may vary with city size: large cities may exhibit stronger within-sector outperformance due to density economies and higher-quality firms, while smaller cities may display more pronounced within-sector convergence (catch-up) dynamics.

In parallel, it is not hard to show that total wage growth risk over the same horizon can be approximated by an analogous decomposition into three components:

$$\sigma^2_{c,T} \approx \underbrace{s'_{c,0} \mathbf{\Omega}^N_T s_{c,0}}_{\text{Benchmark risk exposure}} + \underbrace{\left(\overline{s'_{c,t}\, \mathbf{\Omega}^N_T\, s_{c,t}} - s'_{c,0}\, \mathbf{\Omega}^N_T\, s_{c,0} \right)}_{\text{Reallocation risk}} + \underbrace{\overline{s'_{c,t} (\mathbf{\Omega}_{c,T} - \mathbf{\Omega}^N_T) s_{c,t}}}_{\text{Idiosyncratic risk}}, \tag{4}$$

where the first term, $s'_{c,0} \mathbf{\Omega}^N_T s_{c,0}$, captures *benchmark risk exposure*, defined as the variance of real wage growth a city would face if it retained its initial industry employment shares $s_{c,0}$ throughout the horizon and were exposed to the national variance–covariance structure of industry wage growth, $\mathbf{\Omega}^N_T$. This component reflects passive risk exposure inherited from initial industrial composition: for example, cities initially specialized in highly volatile or strongly co-moving industries—such as energy-oriented cities like Houston or manufacturing-intensive cities like Detroit—mechanically exhibit higher benchmark risk even absent any local adjustment. The second term captures *reallocation-driven risk changes* and is defined as the difference between the time-averaged quadratic form $\overline{s'_{c,t} \mathbf{\Omega}^N_T s_{c,t}}$ and the benchmark exposure $s'_{c,0} \mathbf{\Omega}^N_T s_{c,0}$. Holding the national covariance matrix fixed isolates the contribution of changes in industry weights, highlighting how cities actively mitigate or amplify wage risk through reallocation: cities able to shift employment away from negatively shocked, highly volatile, or strongly co-moving industries toward more stable or weakly correlated

[2] Notable, throughout the paper, "idiosyncratic" growth and risk are defined relative to national industry benchmarks, rather than as micro-level noise that can be diversified away.

sectors—such as diversified service-oriented metros—can reduce overall risk, whereas cities facing high adjustment frictions remain exposed to adverse risk structures. The third term captures *idiosyncratic risk*, defined as the time-averaged quadratic form $\overline{s'_{c,t}(\mathbf{\Omega}_{c,T} - \mathbf{\Omega}_T^N)s_{c,t}}$, which reflects city-specific deviations in the variance–covariance structure of industry wage growth. This component captures how local production networks and shock-transmission mechanisms—often more constrained in smaller or less diversified cities—shape realized wage risk beyond what is implied by national industry risk alone.

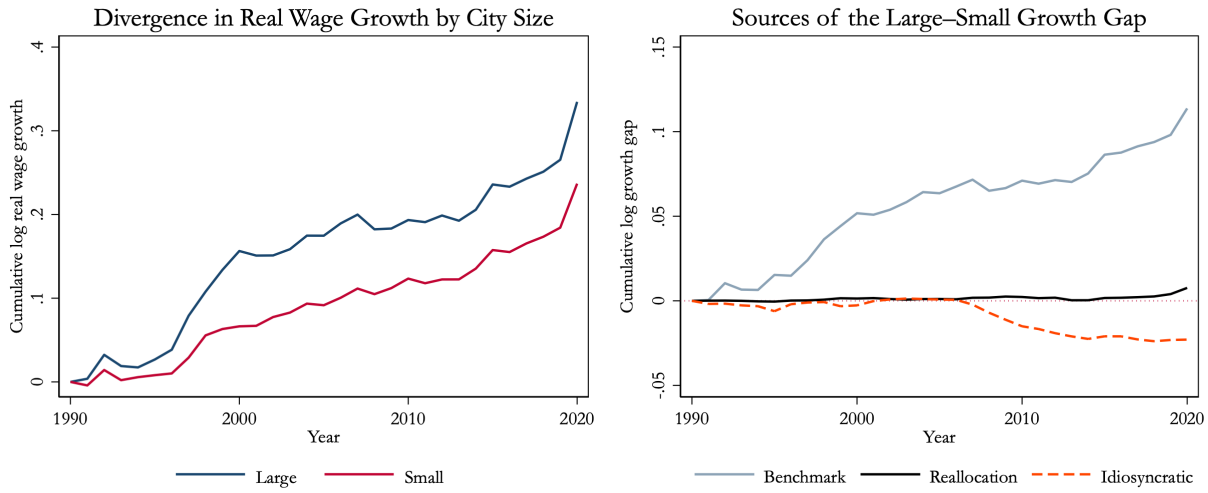
Under this decomposition, we now turn to the empirical drivers of growth and risk efficiency in large versus small cities over the 1990–2020 period. Figure 3 summarizes the evidence on real wage growth. Panel (a) reports cumulative real wage growth for large cities (CBSAs with employment above 0.1 million) and small cities (CBSAs with employment below 0.1 million), documenting a modest growth advantage associated with city size. Panel (b) decomposes this advantage by plotting the large–small difference in each of the three components of Eq. 2.3.1. The figure shows that large cities outperform primarily through higher benchmark growth—reflecting a more favourable initial industry mix, but not through a larger idiosyncratic growth component or reallocation.[3] Panel (c) plots the cumulative contribution of each component separately for large and small cities, showing that benchmark growth associated with initial industry mix accounts for the largest share of overall city growth, echoing evidence from finance that benchmark portfolio weights explain most realized returns (on average 93.6% in Brinson et al. (1986)). Taken together, Figure 3 indicates that benchmark growth is the dominant component of city wage growth and accounts for most of the large-city growth advantage.

We treat the second moment of wage growth as a margin of importance comparable to mean wage growth. Figure 4 summarizes the three components of risk in Eq. 2.3.1 and delivers two key insights. First, large cities are not ex ante safer under the national industry variance–covariance matrix: across the three decades, they exhibit higher benchmark risk exposure, indicating greater exposure to systematic covariance risk embedded in the national industry structure. Second, large cities nonetheless display lower realized wage volatility ex post, mainly because they have lower idiosyncratic risk and, to a lesser extent, lower reallocation risk. This implies that the risk advantage of large cities is driven primarily not by lower benchmark exposure, but by smaller city-specific departures from the national variance–covariance benchmark in the industry pairs most relevant to their employment portfolios.

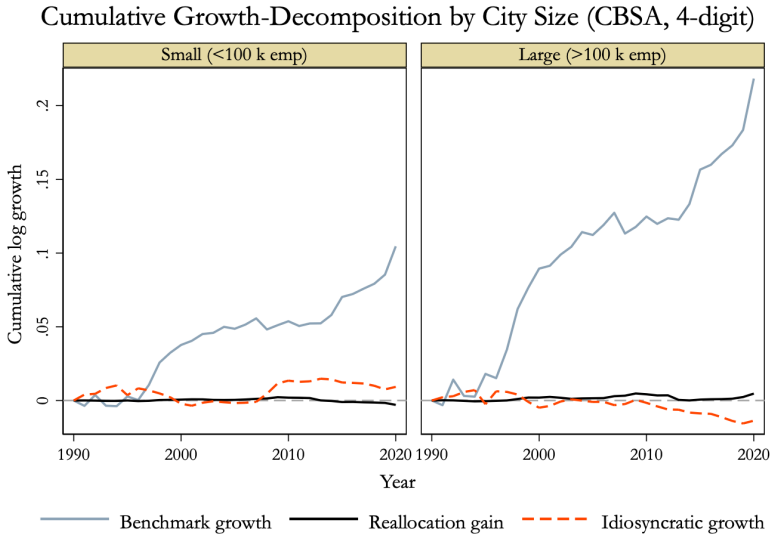
2.3.2 Discussion of mechanisms

Given that lower idiosyncratic risk—defined as the contribution of city-specific deviations in the variance–covariance structure of industry wage growth from the national benchmark—is the primary driver of large cities' relative risk efficiency, it is natural to ask how large cities "manage risk" beyond their inherited exposure under initial industry composition. Figure 5 further decomposes this idiosyncratic risk component into overall contributions from industries' own variances and from pairwise covariances across industries. Formally, the idiosyncratic risk term in Eq. 2.3.1 can be

[3] Lower idiosyncratic growth component may appear counter-intuitive, as one might expect stronger within-industry growth in large cities due to superior firm quality and density economies. We reconcile this finding with inter-city growth convergence documented in Ganong and Shoag (2017). Indeed, conditional on initial wage levels, the idiosyncratic growth component is higher in large cities.



(a) (b)



(c)

Figure 3: Decomposition analysis of real wage growth. Panel (a) plots cumulative real wage growth for large and small CBSAs over 1990–2020. Panel (b) shows the cumulative large–small gap in each component from Eq. 2.3.1. Panel (c) reports the cumulative contributions of benchmark, reallocation, and idiosyncratic growth by city size. Real wage growth is constructed from QCEW data. All series are unweighted city averages, reflecting the behavior of the typical city rather than population-weighted aggregates. Large (small) cities are CBSAs with initial (1990) employment above (below) 0.1 million.

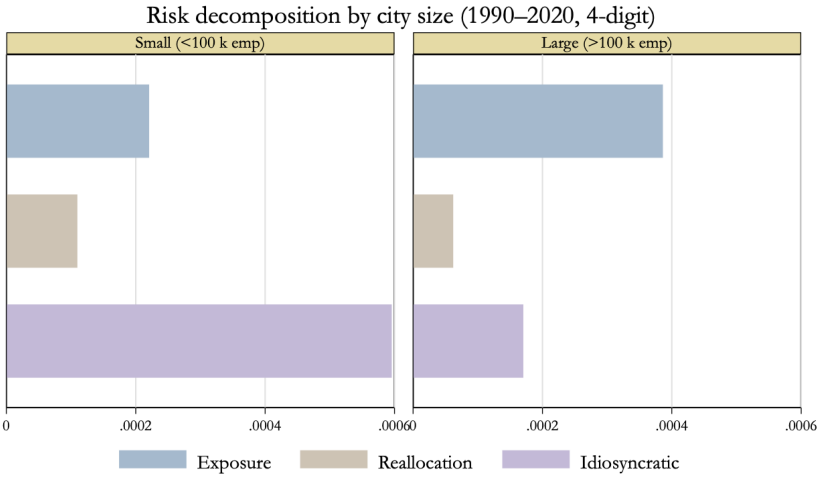


Figure 4: Risk decomposition of real wage growth risk, measured as the time-series variance of real wage growth by city size. The figure reports the benchmark, reallocation, and idiosyncratic risk components defined in Eq. 2.3.1. Bars show unweighted averages across CBSAs over 1990-2020, plotted separately for large and small cities. Data are from the QCEW.

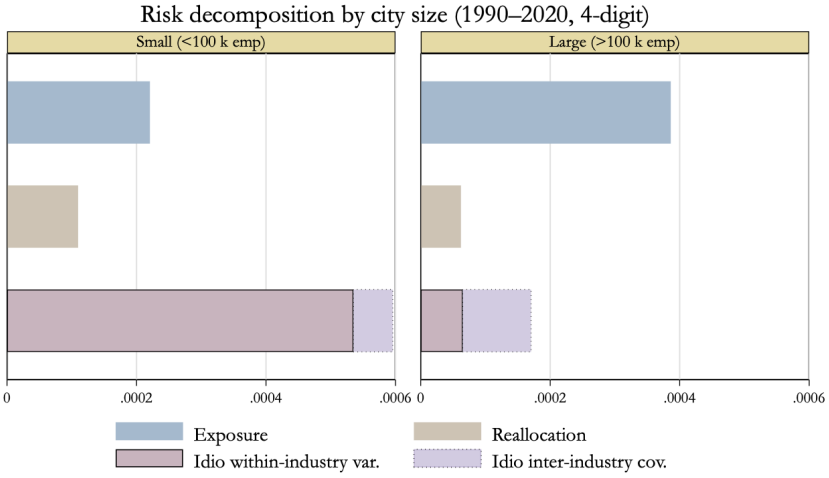


Figure 5: Decomposition of idiosyncratic wage-growth risk into within-industry and inter-industry components by city size. The figure reports the idiosyncratic risk term in Eq. 2.3.1, further decomposed into idiosyncratic within-industry variance (diagonal elements) and idiosyncratic inter-industry covariance (off-diagonal elements). Bars show unweighted averages across CBSAs over 1990-2020, plotted separately for large and small cities. Data are from the QCEW.

written as:

$$\overline{s'_{c,t}(\mathbf{\Omega}_{c,T} - \mathbf{\Omega}_T^N)s_{c,t}} = \underbrace{\overline{\sum_i s_{i,c,t}^2\big(\sigma_{i,c,T}^2 - \sigma_{i,T}^{2,N}\big)}}_{\text{Idiosyncratic variance}} + \underbrace{\overline{\sum_{i\neq j} s_{i,c,t}s_{j,c,t}\big(\sigma_{ij,c,T} - \sigma_{ij,T}^N\big)}}_{\text{Idiosyncratic covariance}}, \tag{5}$$

where $s_{i,c,t}$ denotes the employment share of industry i in city c at time t, $\sigma_{i,c,T}^2$ and $\sigma_{ij,c,T}$ denote, respectively, the local variance of industry i and the covariance between industries i and j over horizon T, and superscript N indicates their national counterparts. The first term therefore captures within-industry variance deviations (idiosyncratic variance), while the second captures inter-industry covariance deviations (idiosyncratic covariance). A key observation from Figure 5 is that the lower idiosyncratic risk observed in large cities is driven primarily by systematically lower within-industry variance, rather than by weaker inter-industry covariance. This pattern suggests that large cities exhibit lower risk primarily through more stable within-industry outcomes, rather than through a reshaping of cross-industry co-movement.[4]

To understand the relatively small contribution of within-industry variance in Eq. 2.3.2, it is useful to recall a basic diversification logic from portfolio theory. As industry employment shares become more dispersed, the contribution of assets' own variance mechanically declines with diversification: as shares spread across industries, the term $\sum_i s_{i,c,t}^2$ decreases, reducing the weight placed on industry-specific variance and shifting importance toward covariance terms. Importantly, diversification is not the only force shaping overall contribution of within-industry risk. Even holding the degree of diversification fixed, cities may differ in their exposure to idiosyncratic variance if their constituent industries are, on average, more volatile relative to national benchmarks. To separate these two channels, we rewrite the within-industry idiosyncratic risk component as the product of an HHI-based measure of industry concentration and a concentration-adjusted average excess industry variance. This decomposition allows us to distinguish the mechanical effect of diversification from differences in the volatility profile of industries that shapes overall within-industry risk:

$$\overline{\sum_i s_{i,c,t}^2\big(\sigma_{i,c,T}^2 - \sigma_{i,T}^{2,N}\big)} = \underbrace{\overline{\sum_i s_{i,c,t}^2}}_{\text{HHI}} \times \underbrace{\frac{\overline{\sum_i s_{i,c,t}^2\big(\sigma_{i,c,T}^2 - \sigma_{i,T}^{2,N}\big)}}{\overline{\sum_i s_{i,c,t}^2}}}_{\text{HHI-adjusted excess variance}}. \tag{6}$$

and Figure 6 plots the two components in Eq. 6. Large cities exhibit both greater industry diversity (lower HHI) and lower HHI-adjusted excess variance. Even conditional on diversity level—which accounts for most of the variation (a slope of -0.025)—constituent industries in larger cities on average are less volatile relative to national benchmarks (a slope of 0.001). This echoes the labour-market pooling view that thicker local labour markets allow establishments within a sector to better absorb idiosyncratic establishment shocks (Krugman, 1991; Overman and Puga, 2010). More broadly, the within-industry risk advantage might operate through two channels: a larger number of establishments, which generates diversification of establishment-level shocks within sectors, and higher establishment quality, which dampens the pass-through of sectoral shocks to wages through richer internal adjustment margins such as reorganization and product-mix adjustment. The latter hypoth-

[4]Appendix B.2 provides additional evidence based on city–industry–level regressions. The results show that, conditional on industry, within-industry wage-growth variance is on average lower in larger cities, whereas, conditional on an industry pair, cross-industry covariance is on average higher in larger cities. This pattern is consistent with stronger inter-industry linkages and home-market effects in large cities, which tend to raise co-movement across related sectors.

esized channel aligns with recent work showing that larger and less distorted firms optimally invest more in risk management, as the gains from reducing productivity risk scale with firm size whereas parts of the organizational and risk-management costs are largely size-invariant, leading to lower productivity volatility and weaker comovement with aggregate conditions (Kopytov et al., 2025). Using Spanish data, (Kopytov et al., 2025) show that larger firms are more likely to have Chief Risk Officer and implement Enterprise-wide Risk Management systems. Figure 7 documents a negative relationship between the HHI-adjusted excess variance and both the number of establishments (panel a) and establishment quality—proxied by the city employment share in medium-to-large establishments with more than 100 employees (panel b), consistent with the relevance of both channels.

Taken together, the decomposition shows that benchmark exposure largely governs the level of city growth and risk through path dependence embedded in initial industry composition. At the same time, large cities distinguish themselves through superior active adjustment, reflected in stronger reallocation gains in wage growth and systematically lower idiosyncratic components of wage-growth volatility. These two margins map naturally into distinct underlying mechanisms. Reallocation gains hinge on cities' ability to redeploy labour across industries, a process constrained by the across-industry portability of workers' skills, which determines how easily displaced workers can move from declining and volatile industries to more stable and expanding ones without large productivity losses or prolonged unemployment spells, thereby shaping both average wage growth and the volatility of wage adjustments. By contrast, the idiosyncratic risk advantage of large cities reflects weaker local amplification of industry-specific shocks, operating primarily through more stable within-industry outcomes rather than reduced cross-industry co-movement, and is plausibly linked to establishment-level characteristics such as firm scale and internal adjustment capacity.

While the preceding analysis provides a transparent accounting of the sources of large cities' growth and risk advantages, it remains descriptive and does not establish causality. Initial industry mix and across-industry diversification are slow-moving state variables that shape cities' long-run exposure to industry-level risk, but do not by themselves govern how shocks are transmitted once they occur. Similarly, within-industry establishment counts primarily affect exposure to firm-level idiosyncratic risk, rather than the propagation of shocks that originate at the industry level. To isolate the causal role of local adjustment capacity in shaping shock transmission, Section 4 exploits the China trade liberalization as an exogenous disturbance to local labour markets. We focus on two margins that plausibly govern the propagation of industry shocks—skill portability and establishment quality—and examine whether cities with greater average worker skill portability and a higher share of employment in medium-to-large establishments exhibit a more muted transmission of industry shocks into wage fluctuations, using an event-study framework with heterogeneous shock intensity across cities arising from differences in initial industrial structure. Appendix B.3 presents a heuristic theoretical framework that helps rationalize the intuition behind these two mechanisms.

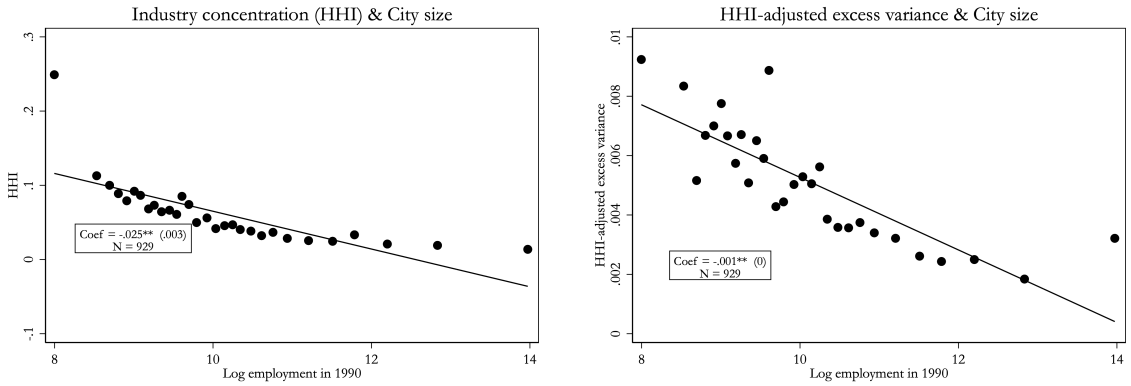


Figure 6: Industry concentration and HHI-adjusted contribution of within-industry idiosyncratic risk by city size. Panel (a) presents a binscatter of cities' initial industry concentration, measured by the Herfindahl–Hirschman Index (HHI) in 1990, against log city employment. Panel (b) presents a binscatter of the HHI-adjusted excess within-industry variance component defined in Eq. 6, plotted against log city employment and averaged across three decades. In both panels, the solid line shows fitted values from an OLS regression, with the estimated slope reported in parentheses. Data are from the QCEW.

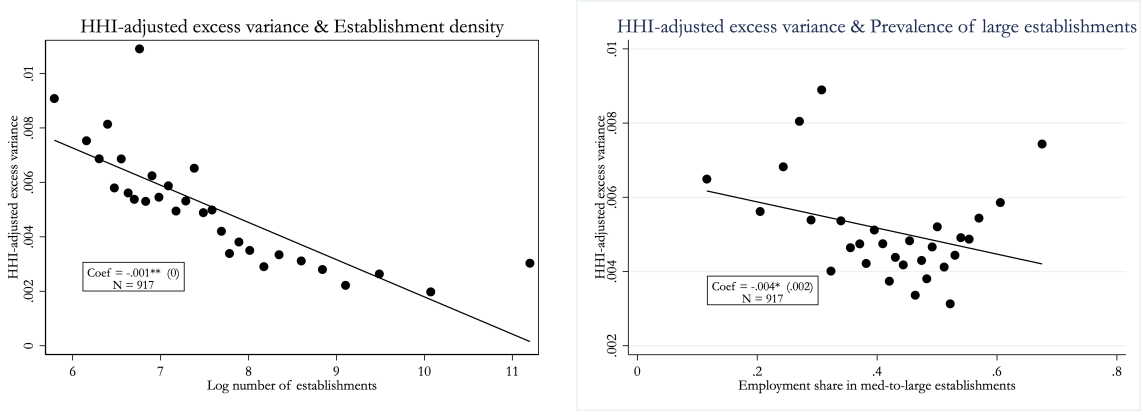


Figure 7: HHI-adjusted excess within-industry idiosyncratic risk and city-level establishment characteristics. Panel (a) presents a binscatter of the HHI-adjusted excess within-industry variance component defined in Eq. 6 against the log number of establishments. Panel (b) presents a binscatter of the same measure against establishment quality, proxied by the city employment share in medium-to-large establishments with more than 100 employees. In both panels, the solid line reports fitted values from an OLS regression, with the estimated slope reported in parentheses. Data are from the QCEW and the CBP.

3 DATA

This section describes the data used in the analysis. It first introduces the wage and employment data used to construct city-level measures of wage growth and time-series wage risk underlying the portfolio analogy and growth–risk decomposition. It then presents the data used to construct two city characteristics capturing local adjustment capacity: a novel measure of skill portability and the establishment quality. Finally, it describes the data used in the event-study analysis of the China trade liberalization shock, which illustrates the transmission of sectoral shocks across cities. See Appendix C for more details.

3.1 Wage and employment

The data on annual wages and employment counts from 1990 to 2020 are drawn from the publicly-available Bureau of labour Statistics Quarterly Census of Employment and Wages (BLS QCEW), which provides quarterly county–industry administrative records based on employer filings under state unemployment insurance (UI) programs and is aggregated to 2013 CBSAs. This paper focuses on the private sector, where wages and employment are primarily determined by market forces rather than institutional pay-setting. Our baseline analysis uses NAICS 4-digit industries, comprising 305 unique industries observed over the sample period. Wages are measured as QCEW average annual pay, defined as total annual payroll divided by average monthly employment, which provides a comprehensive, administratively recorded measure of per-worker labour income, including salaries, bonuses, tips, and other cash compensation, though it may also reflect changes in hours and workforce composition. Quarterly data are averaged to annual values and deflated using the BLS CPI-U. Real wage growth is computed as logarithm difference $g_{cit} = \ln w_{cit} - \ln w_{ci,t-1}$, and real wage risk for a given period is measured as $\sigma_{cip} = \mathrm{sd}(\{g_{cit}\}_{t \in p})$.

3.2 Skill portability

Skill portability is treated as an initial city characteristic measured prior to the China trade shock. We construct skill portability using individual-level data from the 1990 5% IPUMS sample. The sample is restricted to employed, private-sector wage and salary workers who worked at least 40 weeks in the previous year and usually worked at least 35 hours per week, and their industry and occupation information is used to construct baseline city-level measures. [5]

A growing literature emphasizes that human capital specificity is primarily skill- rather than industry-based (Poletaev and Robinson, 2008), implying that cities' adjustment to sectoral shocks depends on the extent to which existing skills can be redeployed across industries. Most existing measures of skill portability focus on individual outcomes, such as task similarity or wage changes following job switches (Gathmann and Schönberg, 2010; Yi et al., 2017; Manning and Aguirre, 2025). While informative about individual mobility, these approaches do not capture cross-city differences

[5] Because the 1990 state sample identifies geography only at the PUMA level and suppresses metropolitan area information, we follow Albouy et al. (2016) and assign counties and CBSAs to PUMAs using the MABLE98/Geocorr1990 crosswalk provided by the Missouri Census Data Center (`https://mcdc.missouri.edu/applications/geocorr1990.html`).

in the aggregate redeployability of skills or cities' capacity to reallocate workers across industries in response to large sectoral shocks.

To capture this spatial difference in labour mobility, this paper develops a novel measure of a city's *skill portability* as the expected probability that a worker can reallocate to another local industry while remaining in the same occupation.[6] This measure builds on the assumption that each occupation represents a bundle of occupation-specific skills that can be deployed across multiple local industries. Consider a worker employed in occupation–industry cell (o, j) in city c. Conditional on staying in occupation o, the probability of being re-employed in a different industry $j' \neq j$ equals $(1 - s^c_{j|o})$, where $s^c_{j|o} \equiv E^c_{oj}/E^c_o$ denotes the employment share of industry j within occupation o in city c. Aggregating over all occupation–industry cells yields the city-level average reallocation probability:

$$P_c = \sum_{o,j} s^c_{oj}\big(1 - s^c_{j|o}\big) = \sum_o s^c_o D^c_o, \qquad D^c_o = 1 - \sum_j \big(s^c_{j|o}\big)^2, \tag{7}$$

where $s^c_{oj} = E^c_{oj}/L^c$ is the fraction of total city employment L^c accounted for by occupation–industry cell (o, j), and $s^c_o = E^c_o/L^c = \sum_j E^c_{oj}/L^c$ denotes the share of occupation o in the city's overall employment. The first equality follows from the law of total probability: the overall reallocation probability equals the weighted average of within-occupation conditional probabilities. Collecting terms by occupation and substituting $s^c_{j|o} = E^c_{oj}/E^c_o$ yields the second equality, where skill portability is expressed as an occupation-share-weighted measure of local cross-industry dispersion, $D^c_o = 1 - \sum_j \big(s^c_{j|o}\big)^2$, which increases when a given occupation is more evenly distributed across industries.

Intuitively, P_c captures a city's reallocation flexibility—the likelihood that workers can shift across industries without changing occupations when an industry-specific shock occurs. Occupations concentrated in a single industry (e.g., *airline pilots*, *steelworkers*) rely on highly specific skills and contribute little to P_c, whereas occupations spanning multiple sectors (e.g., *accountants*, *software engineers*, *sales managers*) embody more general, portable skills and contribute more. By construction, P_c is related to a city's industrial diversity but emphasizes its occupational and human-capital dimension, capturing how broadly each occupation's skills are deployed across industries. As a probability-based measure, P_c also reflects the scale advantage of larger cities, which tend to host a more complete set of occupations within each industry. Cities with higher P_c should therefore be more resilient to sectoral shocks, adjusting more efficiently through within-occupation reallocation and experiencing lower uncompensated wage volatility. As shown in Fig. 8, skill portability lies predominantly between 0.55 and 0.8 and is higher in larger cities, reflecting their broader occupational presence across industries, but increases with city size at a diminishing rate as occupation–industry

[6] A closely related measure of skill transferability is proposed by Yi et al. (2024), who infer cross-sector economic distance from wage changes of exogenously displaced workers who switch sectors. While their empirical design carefully addresses selection into switching, the resulting wage-based measure captures the *price* of reallocation conditional on switching. By contrast, our probability-based measure of skill portability captures the *feasibility* of within-occupation reallocation implied by local industry structure. Moreover, wage changes upon switching may still reflect contemporaneous demand conditions in destination sectors and therefore need not map one-to-one into underlying skill compatibility. The trade-off of our approach is that it abstracts from wage outcomes and treats human capital as predominantly occupation-specific, an assumption supported by evidence that occupations account for substantially more variation in worker outcomes than sectors (e.g. Kambourov and Manovskii (2009); Traiberman (2019)). Nevertheless, other strands of the literature emphasize the role of industry-specific human capital (e.g. Neal (1995); Parent (2000)).

coverage approaches completeness.

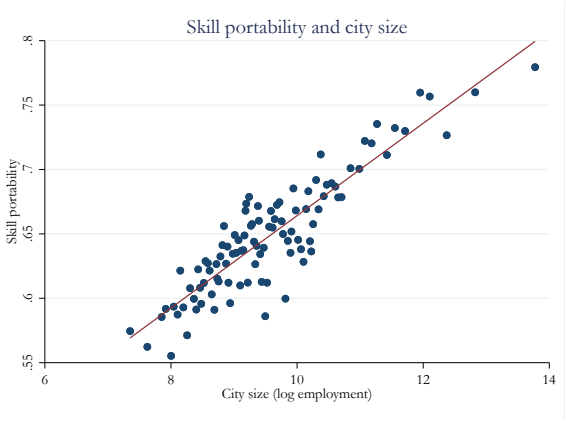


Figure 8: Skill portability and city size. Skill portability is constructed from the 1990 5% IPUMS sample as the occupation-share-weighted probability of reallocating across local industries while remaining in the same occupation (see equation 7). Cities (CBSAs) are grouped into 100 quantiles by log employment; points show within-bin means.

3.3 Establishments and leading-firm

The U.S. County Business Patterns (CBP) dataset provides county–industry information on the number of establishments by employment size class.[7] Using the 1990 CBP data, this paper constructs city-level measures of total establishments and the prevalence of medium-to-large establishments. The latter is measured as the share of city employment accounted for by establishments with more than 100 employees and is used as a proxy for the prevalence of larger production units with presumably greater organizational and adjustment capacity. Because establishment-level employment is not directly observed, employment for each establishment is approximated using the midpoint of its employment size class (e.g., 375 for the 250–499 category).

While the prevalence of large establishments captures one dimension of a city's production structure, it does not directly reflect firms' financial strength and organizational capacity. We therefore complement the analysis with firm-level information from Compustat North America, which provides annual financial data for publicly traded U.S. and Canadian firms since 1950. We identify each firm's location using the CBSA of its headquarters and use data from the 1990 fiscal year (Fundamentals Annual and Historical Segments) to construct pre-determined measures of corporate presence and characteristics at the city level. First, we construct a measure of historical leading-firm presence in each CBSA: the share of publicly traded firms relative to the total number of establishments in CBP. Second, we compute employment-weighted city-level measures summarizing the financial and operational characteristics of publicly traded firms headquartered in each CBSA. These include indicators of financial resilience (leverage, interest coverage ratio and cash-to-assets ratio), operational efficiency (return on assets, operating margin and asset turnover), investment capacity (capital expenditure intensity and R&D intensity), asset structure (PPE-to-assets), and business diversification (business-segment count, industry breadth and Herfindahl index of segment-level sales).

[7] Publicly available at `https://www.census.gov/programs-surveys/cbp/data/datasets.html`.

In the baseline analysis, we include all publicly traded firms, allowing city-level buffering capacity to reflect overall corporate strength. As a robustness check, we restrict the sample to manufacturing firms (SIC 2000–3999) to test whether the effects are driven by within-manufacturing heterogeneity. All measures are constructed using 1990 data to ensure they are predetermined with respect to subsequent trade exposure. Although Compustat primarily covers publicly traded firms and therefore reflects the right tail of the firm-size distribution, listed firms tend to account for a disproportionate share of employment, capital allocation, and strategic decision-making. Their financial and organizational characteristics may therefore shape how trade shocks are transmitted and absorbed at the city level. Appendix C.2 provides additional information on the construction of these variables and the limitations of the Compustat data.

3.4 PNTR exposure and related controls

The PNTR exposure data are drawn from Pierce and Schott (2020), and the empirical implementation follows Pierce and Schott (2016, 2020); Pierce et al. (2024). Exposure to the granting of permanent normal trade relations (PNTR) to China is measured using industry-level *NTR gaps*, defined as the difference between non-NTR tariff rates and the NTR rates applied to WTO members for each SIC industry j:

$$NTRGap_j = NonNTRRate_j - NTRRate_j.$$

Prior to 2000, China's access to NTR tariffs was temporary and subject to annual congressional renewal, creating policy uncertainty for US–China trade. The passage of PNTR in 2000, ahead of China's accession to the WTO in 2001, eliminated this uncertainty by permanently locking in low tariff rates and generated an industry-specific trade liberalization concentrated in manufacturing. Industries with larger NTR gaps therefore experienced larger reductions in trade-policy risk following PNTR and greater subsequent exposure to the China trade shock. As shown by Pierce and Schott (2016), most cross-industry variation in the NTR gap is driven by historically predetermined non-NTR tariff rates, leaving little scope for reverse causality from contemporaneous industry conditions. CBSA-level PNTR exposure is constructed by aggregating industry-level NTR gaps using pre-policy (1990) employment shares:

$$PNTRexp_c = \sum_j \frac{L_{jc}^{1990}}{L_c^{1990}} NTRGap_j,$$

where c indexes CBSAs, j indexes industries, and L_{jc}^{1990}/L_c^{1990} denotes the share of CBSA c's total employment in industry j in 1990. As shown in Figure 9, the distribution of $PNTRexp_c$ has a mean of 7.15 percent and a standard deviation of 4.98 percent, with an interquartile range of 6.29 percentage points.

Following Pierce and Schott (2020), the baseline specification controls for contemporaneous policy changes and demographic characteristics. Two time-varying controls are included. First, we control for the average U.S. import tariff faced by each CBSA, constructed as the employment-share-weighted tariff across its active industries. Second, we account for local exposure to the removal of quantitative restrictions on textile and apparel imports associated with the phase-out of the Multi-Fiber Arrangement (MFA).

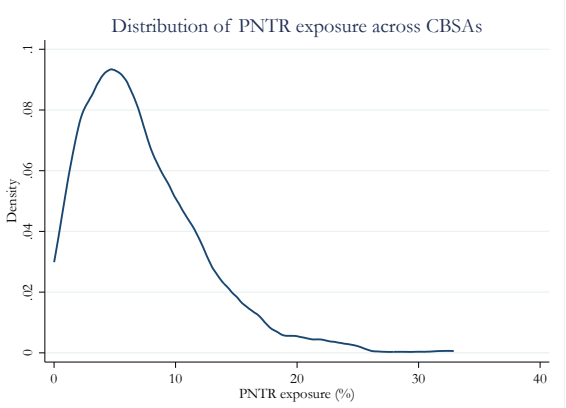


Figure 9: Kernel density of PNTR exposure across CBSAs. PNTR exposure is constructed as the employment-share-weighted average of industry-level NTR gaps using 1990 employment shares. Data is based on Pierce and Schott (2020).

We allow for time-varying effects of several CBSA characteristics fixed at their 1990 levels to isolate the causal impact of the PNTR shock: (i) the homeownership rate, controlling for possible time-variant housing lock-in effects and mobility frictions. (ii) the initial unemployment rate, proxying potential national time-variant shocks to persistent slack labour market; (iii) log average household income, to control for wealth-based resilience and baseline affluence; and (iv) the share of foreign-born population, to account for immigration-related labour supply shifts. Finally, we control for cities' exposure to changes in Chinese trade policy following WTO accession, measured by Chinese import tariffs imposed on US goods and domestic production subsidies.

4 CHINA TRADE SHOCK AND WAGE DYNAMICS

4.1 Background and Event-Study Design

As the decomposition analysis in Section 2.3 shows, divergence in cities' long-run portfolio performance arises not only from initial portfolio structure (i.e., industry weights, industry diversity, number of firms), but also from rebalancing and risk attenuation capacity. However, urban wage growth dynamics and cities' relative positions with respect to the efficient frontier reflect the accumulation of many shocks across industries, making it difficult to isolate the city characteristics that govern growth and risk using standard OLS regressions. For example, the apparent "efficiency" of large cities may be confounded by unobserved fundamentals—such as amenities or institutions—that simultaneously affect wages and city size through migration and firm sorting. We therefore exploit the China trade shock induced by U.S. trade liberalization as an exogenous negative labour demand shock that varies with cities' initial industry composition, and implement an event-study design to identify whether pre-determined city characteristics related to adjustment capacity—such as skill portability and firm quality—shape the transmission and absorption of the shock. More specifically, we follow Pierce and Schott (2020) and use the U.S. granting of Permanent Normal Trade Relations (PNTR) to China as a shock with a well-defined timing, which enables an event-study analysis of dynamic wage responses. This contrasts with long-difference approaches such as Autor

et al. (2013), which relate changes in import exposure across Commuting Zones (CZs) to decadal changes in local labour market outcomes. Notably, the analysis focuses on a negative shock, as adverse disturbances simultaneously increase wage risk and depress wage growth, making cross-city differences in shock-absorption capacity more salient than under positive shocks and more relevant for overall city efficiency.

The China trade shock. Before 2000, China's access to U.S. Normal Trade Relations (NTR) tariff rates was temporary and subject to annual congressional renewal, generating substantial policy uncertainty in U.S.–China trade. Failure to renew NTR would have subjected most Chinese goods to the much higher non-NTR tariffs set under the Smoot–Hawley Tariff Act of 1930, creating significant uncertainty for U.S. firms' long-run sourcing from China. The U.S. passage of Permanent Normal Trade Relations (PNTR) in October 2000—prior to China's accession to the WTO in 2001—eliminated this renewal uncertainty by permanently granting China NTR tariff treatment. This institutional change constituted an industry-specific trade liberalization concentrated in manufacturing, such that cities specialized in industries with larger NTR gaps experienced larger declines in trade-policy risk and greater subsequent exposure to the China trade shock. Details on industry-level NTR gaps and city-level PNTR exposure are provided in Section 3.

The abrupt passage of PNTR generated long-lasting, significant, and spatially heterogeneous effects on U.S. employment and wage dynamics through increased exposure to Chinese imports. As shown by Handley and Limão (2017), the reduction in trade-policy uncertainty associated with PNTR accounts for 22%–30% of the post–WTO growth in U.S. imports from China. A growing literature further documents that regions and workers more exposed to PNTR experienced higher unemployment rates and declines in labour force participation following its implementation (Pierce and Schott, 2016, 2020; Pierce et al., 2024; Rodríguez-Clare et al., 2026). Most closely related to our study, Pierce and Schott (2020) show, using matched employer–employee data, that the adverse earnings effects of PNTR operate primarily through workers' local labour market exposure rather than their own-industry exposure, and that indirect exposure through input–output linkages generates heterogeneous effects across both manufacturing and non-manufacturing workers. Together, these findings underscore the importance of spatial exposure and industry network structure, motivating our focus on how city-level characteristics and industry composition shape the propagation of trade shocks and the resulting wage growth and risk dynamics.

Explained variable. We estimate an event-study specification with log wages as the outcome variable. Because log wages accumulate past wage growth,

$$\ln w_{ct} = \ln w_{c,t_0} + \sum_{\tau=t_0+1}^{t} g_{c\tau},$$

the event-study coefficients trace cumulative deviations in wage paths rather than transitory one-period responses. This feature is essential in our setting, as trade liberalization represents a labour demand shock whose effects propagate over time through cities' initial industrial composition. From a portfolio perspective, such adverse shocks jointly reduce growth and increase downside risk, while city characteristics governing reallocation and risk buffering shape how these shocks propagate.

4.2 Identification strategy

This section examines the dynamic effects of PNTR on workers' wages and tests whether city characteristics – skill portability and firm quality – mediate the transmission of this shock exposure with the following dynamic difference-in-differences specification:

$$\begin{aligned}\ln w_{ct} = &\sum_{t\neq 1999} \theta_t \, \mathbf{1}\{year = t\} \times PNTRexp_c \; + \sum_{t\neq 1999} \mathbf{1}\{year = t\} \times PNTRexp_c \times Z_c'\Gamma_t \\ &+ \; X_{ct}\beta \; + \sum_{t\neq 1999} \lambda_t \, \mathbf{1}\{year = t\} \times X_c \; + \; \delta_c + \delta_t + \varepsilon_{ct},\end{aligned} \tag{8}$$

where $\ln w_{ct}$ is the average real wage in city (CBSA) c in year $t \in [1990, 2013] \setminus \{1999\}$, with 1999 serving as the baseline year prior to the policy change; $\mathbf{1}\{year = t\}$ denotes a range of year indicators; $PNTRexp_c$ captures city c's exposure to Chinese import competition induced by the resolution of trade policy uncertainty in 2000, constructed using its 1990 industry employment shares.

The term Z_c represents a vector of pre-determined, time-invariant city characteristics that moderate the impact of the trade shock. To systematically explore the potential shock reduction related to city size, the composition of Z_c evolves across the subsequent empirical analyses. First, to establish the baseline buffering effect of urban scale (Section 4.3), Z_c consists solely of initial city size ($\ln Emp_c$). Second, to evaluate the specific structural mechanisms proposed in the decomposition framework (Section 4.4), Z_c is specified as $(SP_c, EQ_c)'$, where SP_c denotes skill portability and EQ_c denotes firm quality. Finally, to conduct the formal mediation analysis (Section 4.5), Z_c combines both the scale and the mechanism variables, $Z_c = (\ln Emp_c, SP_c, EQ_c)'$.

In this framework, the coefficients θ_t trace the dynamic baseline wage response to PNTR exposure inherent in cities' initial industry composition, echoing the baseline risk exposure in the earlier risk decomposition. Because the moderator variables in Z_c are demeaned in the estimation, θ_t isolates the wage trajectory for an average city in response to the PNTR shock. The parameter vector Γ_t captures the heterogeneity in this response with respect to the included city characteristics, mapping directly to the across-industry and within-industry active adjustment margins identified in the decomposition analysis. Formally, the marginal effect of PNTR exposure on city wages in year t is given by:

$$\frac{\partial \ln w_{ct}}{\partial PNTRexp_c} = \theta_t \; + \; Z_c'\Gamma_t.$$

To isolate the causal role of local adjustment capacities, the specification flexibly controls for a host of potential confounders. X_{ct} includes time-varying controls for tariff exposure (NTR_{ct}) and exposure to the MFA phase-out ($MFAExposure_{ct}$). Furthermore, X_c collects a comprehensive set of initial (1990) city attributes—including the mediator variables, the initial unemployment rate, homeownership rate, average household income, and the share of foreign-born population—as well as exposure to contemporaneous Chinese trade policy (e.g., changes in Chinese tariffs and domestic production subsidies; see Section 3). Allowing these initial characteristics in X_c to have time-varying effects (λ_t) mitigates concerns that unobserved macroeconomic shocks (e.g., differential exposure to the tech boom or state-level policies) might disproportionately affect certain types of cities and be

spuriously attributed to the PNTR response. δ_c and δ_t denote city and year fixed effects respectively. Standard errors are clustered at the CBSA level.

4.3 Baseline Results: City Size and Shock Absorption

This section examines the dynamic wage effects of PNTR exposure and its heterogeneity by initial city size using Equation 8. We set the moderator Z_c to demeaned 1990 log employment ($\ln Emp_c$) to capture the buffering (or amplifying) role of urban scale, while also including $\ln Emp_c$ in the control set X_c to absorb secular size-dependent wage trends. Because Z_c is demeaned, the baseline coefficients θ_t isolate the dynamic PNTR effect for an average-sized city, whereas the interaction coefficients $\Gamma_t = \gamma_{0,t}$ capture the marginal effect of city scale on shock absorption. To transparently convey economic magnitudes, we plot the coefficients scaled by their respective interquartile ranges (IQRs). The plotted baseline effect, $\tilde{\theta}_t = \theta_t \times IQR^{PNTR}$, translates to the implied wage impact of an IQR increase in PNTR exposure on an average-sized city. The plotted interaction effect applies a double-IQR scaling, $\tilde{\gamma}_{0,t} = \gamma_{0,t} \times IQR^{PNTR} \times IQR^{\ln Emp}$, to quantify the relative buffering capacity. It measures the difference in wage trajectories between a 75th- and 25th-percentile city facing an identical IQR shift in trade shock exposure.

Given the large number of estimated parameters, Figure 10 plots the dynamic estimates.[8] To rule out confounding from pre-existing local wage dynamics, we present both a baseline specification (Panels a and b) and a detrended model that partials out city-specific pre-1999 linear trends (Panels c and d).[9] Panels (a) and (c) display the scaled baseline event-study coefficients ($\tilde{\theta}_t$), capturing the wage impact of an interquartile (IQR) shift in PNTR exposure. There is no evidence of differential pre-trends prior to 2000 once city-specific pre-trends are accounted for. Following China's WTO accession, wages fall significantly in more exposed cities. Focusing on the baseline results in Panel (a), this negative effect intensifies through the late 2000s and reaches approximately -0.037 log points by 2010. To interpret this, an IQR increase in PNTR exposure translates to a 3.6% wage decline for an average-sized city ($e^{-0.037} - 1 \approx -3.6\%$).[10]

Panels (b) and (d) report the scaled triple-interaction coefficients ($\tilde{\gamma}_{0,t}$), capturing heterogeneity in the shock's impact by city size. The estimates turn positive and more statistically significant post-2000, suggesting that larger cities attenuated local wage losses. Focusing on the baseline results in Panel (b), this buffering effect reaches approximately 0.01 log points by 2010. To interpret this, moving from the 25th to the 75th percentile of the city size distribution averts a 1% wage drop when facing an identical IQR trade shock ($e^{0.010} - 1 \approx 1.0\%$). Remarkably, this spatial advantage offsets roughly 27% of the baseline wage decline (i.e., 0.01/0.037).[11]

Overall, the negative wage effects of geographic-level China shock exposure align with standard

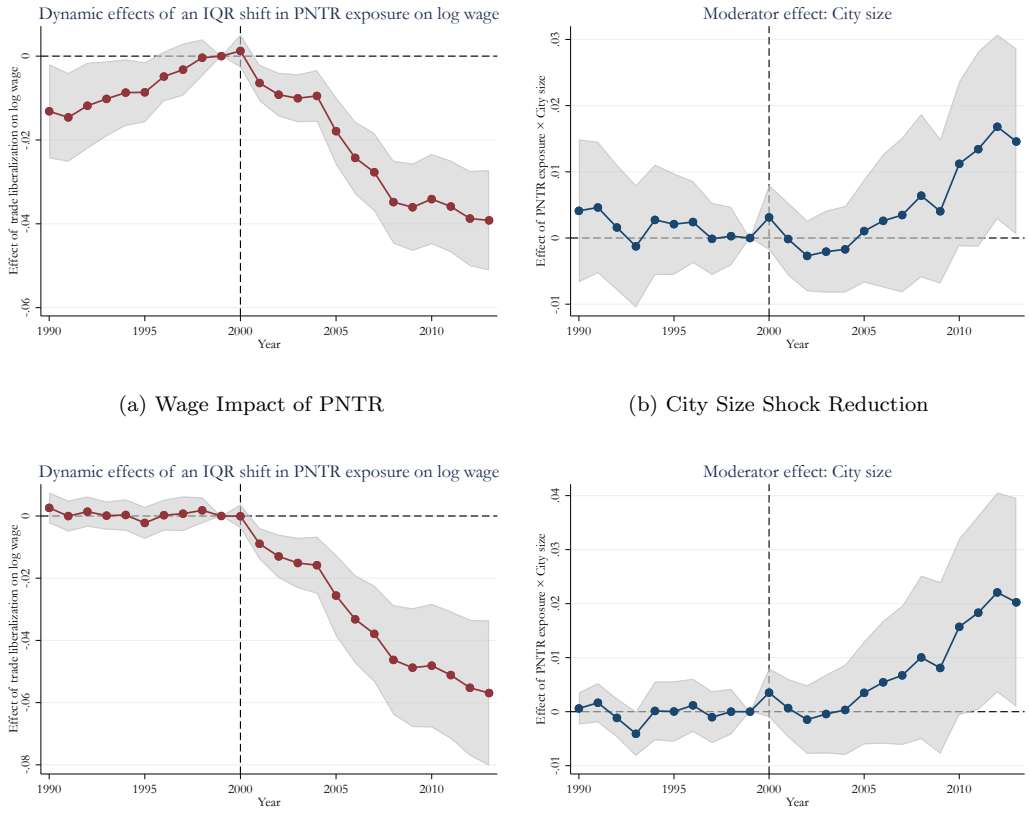
[8] Full regression tables are available upon request.

[9] See Appendix D.1.1 for details of the de-trending procedure. The rest of the main paper focuses on the detrended specification, while results based on the undetrended specification are reported in the Appendix.

[10] Following $\tilde{\theta}_t = \theta_t \times IQR^{PNTR}$, the plotted 2010 value corresponds to the raw point estimate ($\theta_{2010} \approx -0.006$) multiplied by the 6.29 percentage-point IQR in PNTR exposure.

[11] Following $\tilde{\gamma}_{0,t} = \gamma_{0,t} \times IQR^{PNTR} \times IQR^{\ln Emp}$, the plotted 2010 value $\tilde{\gamma}_{0,2010} = 0.010$ corresponds to the raw point estimate ($\gamma_{0,2010} \approx 0.001$) multiplied by the respective IQRs of PNTR exposure (6.29 percentage points) and initial city log employment (1.59 log points). To give an alternative sense of the magnitude, a doubling of city size ($\Delta \ln Emp \approx 0.69$) attenuates the wage effect of an IQR trade shock by about 0.004 log points ($0.001 \times 0.69 \times 6.29 \approx 0.004$), offsetting roughly 11% of the baseline decline.

trade models and empirical evidence. For example, consistent with models where import competition under trade imbalances depresses local labour demand and transmits beyond manufacturing through reduced local expenditure, Autor et al. (2013) find that a $1,000 per worker increase in a Commuting Zone's trade exposure reduced mean weekly earnings by roughly 0.76% over a decade. Using matched employer–employee data, Pierce et al. (2024) document persistent earnings losses for manufacturing workers following PNTR and find that own-county exposure dominates own-industry exposure, even when accounting for input–output linkages. However, relatively little work has examined how local labour-market characteristics beyond industry composition shape how the China shock translates into heterogeneous wage responses. A notable exception is Bloom et al. (2024); while focusing on employment reallocation rather than wages, they document that large high-human-capital coastal labour markets experience net employment gains of about 0.5% per year as service expansion offsets manufacturing losses. Building on this, our results establish city size as a first-order margin of spatial resilience, prompting our subsequent investigation into the mechanisms underlying this scale-based shock absorption.



(a) Wage Impact of PNTR

(b) City Size Shock Reduction

(c) Wage Impact of PNTR (Detrended)

(d) City Size Shock Reduction (Detrended)

Figure 10: Dynamic wage effects of an IQR shift in PNTR exposure and heterogeneity in effects by city size. Panels (a) and (c) plot the scaled baseline event-study coefficients $\tilde{\theta}_t$ capturing the dynamic effect of an interquartile range (IQR) shift in PNTR exposure on city-level log wages (relative to 1999). Panels (b) and (d) plot the scaled triple-interaction coefficients $\tilde{\gamma}_{0,t}$, which capture the added buffering effect of an IQR shift in initial city size (1990 log employment) against the same trade shock. Panels (a) and (b) report estimates from the baseline specification. Panels (c) and (d) report estimates from an alternative specification that partials out city-specific pre-existing linear wage trends (estimated using pre-1999 data). All specifications include city and year fixed effects and flexibly control for time-varying effects of initial city attributes as in Equation 8. Shaded areas denote 95% confidence intervals, with standard errors clustered at the CBSA level. Data from QCEW and Pierce and Schott (2020).

4.4 Mechanisms of Urban Shock Absorption

Having established that large cities buffer the China shock, we now unpack the underlying mechanisms driving this scale advantage. Guided by our decomposition analysis in Section 2.3, we evaluate two specific local adjustment channels: skill portability and firm quality. In terms of the latter, we first use County Business Patterns (CBP) as our full-sample baseline, measuring the prevalence of large local establishments as a proxy for firm quality. However, because CBP lacks financial and organizational information, we supplement it with Compustat data on publicly traded firms to substantiate this quality measure.[12]

4.4.1 Across- and Within-Industry Adjustment Margins

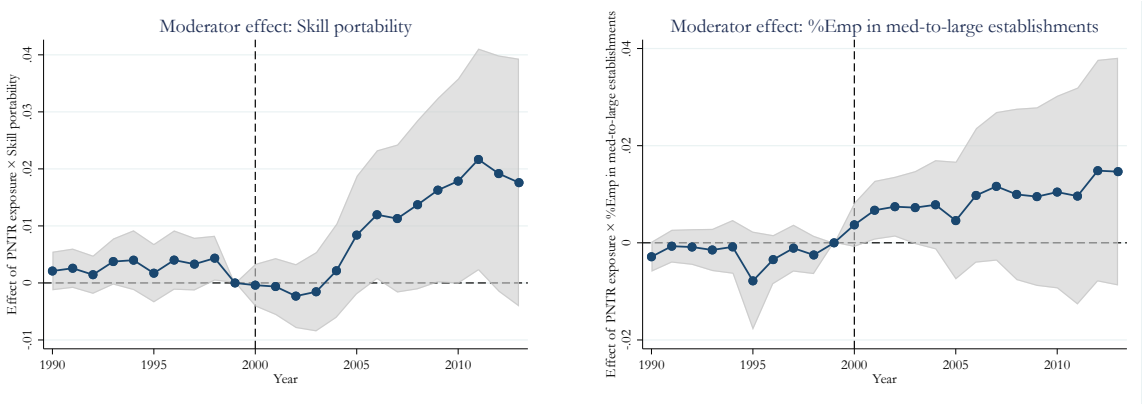
These two mechanisms map directly to the structural adjustment margins identified in our theoretical framework: skill portability captures the across-industry margin by facilitating labour reallocation following sectoral disruptions, whereas establishment quality captures the within-industry margin by mitigating wage volatility through internal organisational buffers. To empirically evaluate these channels, we set the moderator vector $Z_c' = (SP_c,\ EQ_c)'$ in Equation 8, where SP_c denotes the pre-shock (1990) city-level skill portability, defined as the expected probability that a worker can reallocate to another local industry while remaining in the same occupation; and establishment quality, EQ_c, is proxied by city c's initial (1990) employment share in establishments with more than 100 employees. As before, Figure 11 plots the double-IQR scaled coefficients ($\tilde{\gamma}_{1,t}$ and $\tilde{\gamma}_{2,t}$), which directly quantify the wage trajectory difference between a 75th- and 25th-percentile city in the respective mechanism when facing an interquartile shift in PNTR exposure.

Using the detrended wage specification, focusing first on the across-industry margin, Panel (a) reveals that following a remarkably flat pre-trend, the moderating effect of skill portability ($\tilde{\gamma}_{1,t}$) turns positive and statistically significant after 2005. By 2010, this effect reaches approximately 0.018 log points, implying that moving from the 25^{th} to the 75^{th} percentile of the SP distribution averts a 1.8% wage drop ($e^{0.018} - 1 \approx 1.8\%$) and offsets roughly 49% of the baseline aggregate wage decline (i.e., 0.018/0.037). Turning to the within-industry margin in Panel (b), the estimates for establishment quality ($\tilde{\gamma}_{2,t}$) confirm the absence of differential pre-trends and exhibit a positive divergence immediately following China's WTO accession. These effects, however, are only statistically significant in the short run (3 years), likely due to the imprecision of using categorical employment data in the CBP. By 2010, the buffering effect of EQ reaches 0.01 log points, implying that a comparable interquartile improvement in firm quality preserves 1% of local wages, insulating workers from approximately 27% of the baseline shock (i.e., 0.01/0.037). These results demonstrate that both a city's skill portability and its micro-level establishment quality act as significant shock absorbers.

To contextualize these findings, we conceptualize urban shock absorption through both labour-market and firm-level channels. On the across-industry margin, cities with higher skill portability possess a more elastic local labour supply. Portable skills enable displaced workers to reallocate with minimal skill depreciation, thereby averting extreme "cliff-edge" wage collapses and mitigating neg-

[12] While Compustat provides rich data on corporate financial fundamentals, its restriction to publicly listed firms results in zero observations for some smaller micropolitan areas (μSAs). This geographic limitation restricts its use to auxiliary mechanism analysis rather than the full-sample baseline.

ative local demand multiplier effects. This evidence aligns with Bloom et al. (2024), who document more robust post-shock reallocation in high-human-capital labour markets. On the within-industry margin, the buffering role of large establishments is consistent with Kopytov et al. (2025), who identify a scale advantage in corporate risk management, suggesting that larger firms are better equipped to internalize and absorb aggregate shocks. However, while the CBP data indicates that the physical scale of local establishments cushions wage declines, the employment share in medium-to-large establishments remains an indirect proxy for firm quality. To further explore whether this within-industry scale effect genuinely reflects superior financial and strategic resilience, we look beyond physical establishment size—motivating our subsequent analysis of corporate micro-foundations using the Compustat dataset.



(a) Skill Portability Buffering

(b) Med-to-large Establishment Buffering

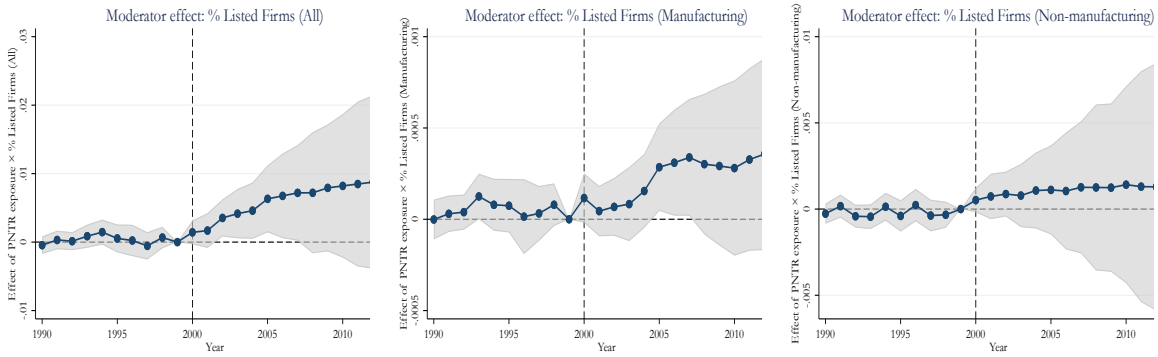
Figure 11: The structural mechanisms of spatial resilience: Skill Portability (SP) and Establishment Quality (EQ). This figure plots the dynamic buffering effects of a city's initial skill network and firm quality against the PNTR wage shock. The estimates are obtained jointly from the baseline event-study specification (Equation 8), allowing SP and EQ to compete in a "horse race." Panel (a) plots the scaled interaction coefficients for Skill Portability ($\tilde{\gamma}_{1,t} = \gamma_{1,t} \times IQR^{PNTR} \times IQR^{SP}$), capturing the differential wage trajectory between a high-SP city (75th percentile) and a low-SP city (25th percentile) confronting an identical IQR trade shock. Panel (b) plots the analogous scaled interaction coefficients for Establishment Quality ($\tilde{\gamma}_{2,t} = \gamma_{2,t} \times IQR^{PNTR} \times IQR^{EQ}$). Both panels report estimates from a specification that partials out city-specific pre-1999 linear wage trends. All specifications include city and year fixed effects, along with flexible controls for time-varying effects of initial city attributes, including city size ($\ln Emp_c$). Shaded areas denote 95% confidence intervals, with standard errors clustered at the CBSA level.

4.4.2 The Micro-Foundations of Firm Quality

To peer into the corporate veil of our CBP-based findings, we lever the Compustat database. Although restricted to publicly traded headquarters—naturally omitting the left tail of the firm-size distribution—these listed companies serve as vital "anchor" employers. Because they account for a disproportionate share of local capital and strategic decision-making, observing their prevalence and characteristics offers a precise window into the organizational capacities that might insulate local labour markets from sectoral shocks.

To operationalize this, we first capture the local prevalence of these corporate anchors by calculating the share of publicly listed firms relative to total city establishments, setting the moderator

variable in Equation 8 to $Z_c = ListedShare_c$.[13] We then separately evaluate the listed firm shares within the manufacturing and non-manufacturing sectors. The results in Figure 12 show that an IQR shift in the total listed share averts a significant 0.6% wage drop by 2005 (Panel a). Sectorally, a high concentration of listed firms strictly within the directly exposed manufacturing sector yields a minor, yet statistically significant, 0.03% buffering effect (Panel b).[14] Conversely, the estimate for non-manufacturing firms remains statistically insignificant (Panel c). This pattern aligns with our earlier decomposition, which attributes spatial resilience to lower *within-industry* volatility. Given that the buffering effect of firm presence is statistically robust only within the directly exposed sector, our subsequent analysis focuses exclusively on the underlying financial characteristics of manufacturing listed firms to unpack this within-industry resilience.



(a) Total Listed Share (b) Manufacturing (c) Non-Manufacturing

Figure 12: The Shock-absorption Effect of Prevalence of Publicly Listed Firms Across Cities. This figure plots the estimated moderating effects of the local share of publicly listed firms on the wage response to the PNTR shock. Panel (a) uses the total share of listed firms relative to all establishments. Panels (b) and (c) decompose this presence into the manufacturing and non-manufacturing sectors, respectively. Shaded areas denote 95% confidence intervals.

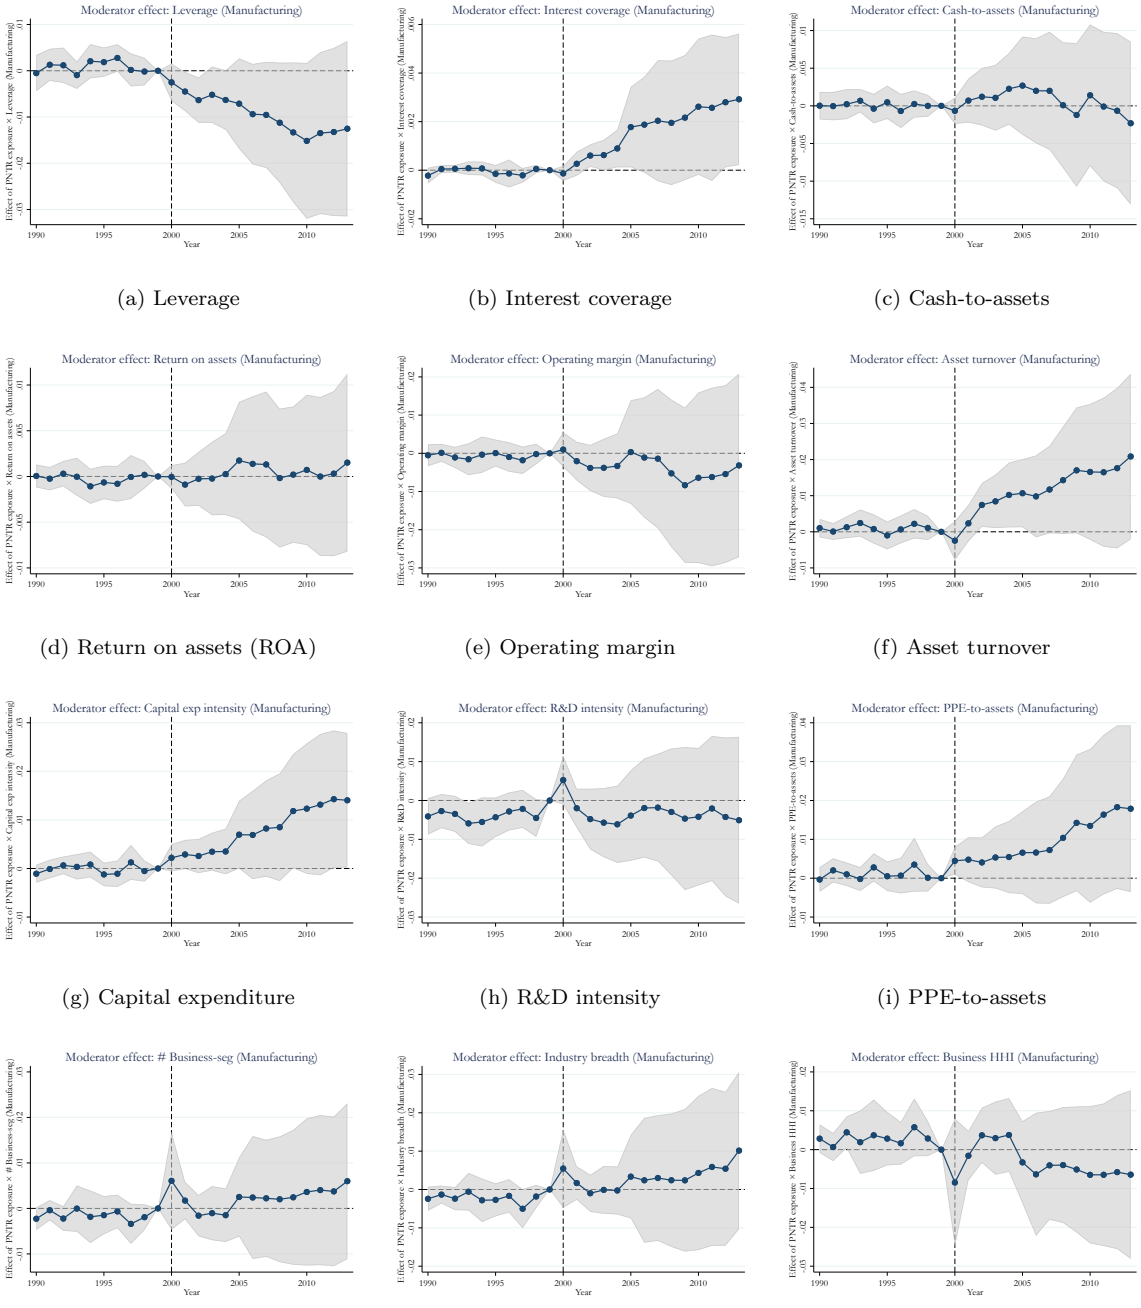
To unpack this within-industry resilience, we examine the intensive margin by evaluating the buffering role of various financial and operational characteristics of local manufacturing listed firms. Specifically, we aggregate firm-level characteristics to the city level, weighted by firm employment. Our results in Figure 13 reveal that at least two dimensions of firm quality serve as critical shock absorbers. First, financial resilience is paramount: cities anchored by firms with lower leverage and higher interest coverage experience significantly smaller wage declines. This suggests that firms with robust balance sheets possess the capacity to absorb revenue shocks without immediate labour shedding or wage cutting. Second, investment and operating efficiency play critical roles; higher capital expenditure intensity and asset turnover facilitate productivity-enhancing upgrades, allowing firms to move up the value chain and escape direct low-end import competition. In contrast, the insignificant estimates for business and industry diversification suggest that the role of internal reallocation seems less important when we focus on listed firms' buffering role.

[13] We rely on the count of listed firms rather than their employment size because Compustat reports consolidated firm-level employment and does not disaggregate the geographic distribution of workers across specific local labour markets.

[14] The substantially larger magnitude of the total effect (0.6%) relative to the manufacturing-specific effect (0.03%) likely reflects general equilibrium dynamics. While manufacturing accounts for only a fraction of local employment, a robust broader corporate ecosystem helps arrest cascading negative demand multipliers (Autor et al., 2013), even though our sector-specific estimate for non-manufacturing firms proves statistically noisy.

Overall, these reveal that the spatial resilience observed in CBP data is fundamentally rooted in the micro-level financial and operational advantages of corporate anchors. Specifically, cities with a dense presence of high-quality listed firms benefit from corporate structures equipped with stronger financial fundamentals—namely, lower leverage and higher interest coverage. These financial buffers are critical for shock absorption, suggesting that large organizations possess superior internal risk-management and operational flexibility (Kopytov et al., 2025). When faced with a severe sectoral shock, these firms can absorb the distress internally through balance sheet strength and productivity-enhancing investments, rather than immediately transmitting the shock to the local labour market through drastic wage cuts or mass lay-offs. Thus, the buffering effect of firm scale is not merely a function of size, but of the qualitative robustness inherent in these corporate structures.

Nevertheless, it is important to note that our findings speak to the mitigation of sectoral shocks, but remain silent on the aggregate risk implications of large establishments. As Gabaix (2011) highlights, when the firm-size distribution is fat-tailed, idiosyncratic shocks to large firms themselves can generate aggregate fluctuations. Thus, while firm scale attenuates the transmission of trade-induced sectoral shocks, it may simultaneously amplify local aggregate risk through reduced diversification of idiosyncratic firm-level shocks.



(a) Leverage (b) Interest coverage (c) Cash-to-assets

(d) Return on assets (ROA) (e) Operating margin (f) Asset turnover

(g) Capital expenditure (h) R&D intensity (i) PPE-to-assets

(j) Business-segment count (k) Industry breadth (l) Segment sales HHI

Figure 13: Dynamic Moderator Effects of Manufacturing Publicly Listed Firm Characteristics. This figure plots the coefficients of the triple interaction between the PNTR shock, various city-level average (employment-weighted) manufacturing firm characteristics, and year dummies. Shaded areas represent 95% confidence intervals based on standard errors clustered at the CBSA level. Data are sourced from Compustat Fundamental Annual and Historical Segment files as of 1990. See Appendix Table C.3 for variable definitions.

4.5 Mediation Analysis

To rigorously evaluate whether the identified channels of skill portability and firm quality constitute the key micro-mechanisms underlying the risk-reduction effect of city size, this section conducts a mediation analysis. This involves comparing the baseline city-size buffering effect—as estimated in Section 4.3—with an augmented specification that jointly controls for the interactions of PNTR exposure with the proposed structural mediators, SP_c and EQ_c, alongside city size $\ln Emp_c$. Specifically, the moderator vector in the event-study specification (Equation 8) is set as $Z_c' = (SP_c,\ EQ_c,\ \ln Emp_c)'$. By allowing these mechanisms to "compete" with urban scale in a horse-race regression, I can identify the extent to which the total scale-dependent resilience is absorbed by these specific labour market attributes.

As illustrated in Figure 14, the event-study coefficient on the city-size triple interaction—which captures the baseline buffering effect of urban scale (depicted by the solid orange line)—exhibits a pronounced decline in magnitude and loses its statistical significance when the model simultaneously accounts for SP_c and EQ_c (depicted by the dashed purple line). Specifically, the inclusion of these mechanisms absorbs a substantial portion of the scale-dependent wage response, with the coefficient on the city-size interaction term decaying by more than 75% over a decade.

This pattern of coefficient attenuation suggests that the industrial-shock risk reduction effect of urban scale is largely mediated through the labour market reallocation flexibility and corporate organizational buffers inherent in larger metropolitan areas. By "opening the black box" of the urban scale effect, these results reinforce the central hypothesis that urban resilience is an endogenous outcome of a city's inherited industry-portfolio structure and its locally mediated adjustment capacity, rather than a mere reflection of population density.

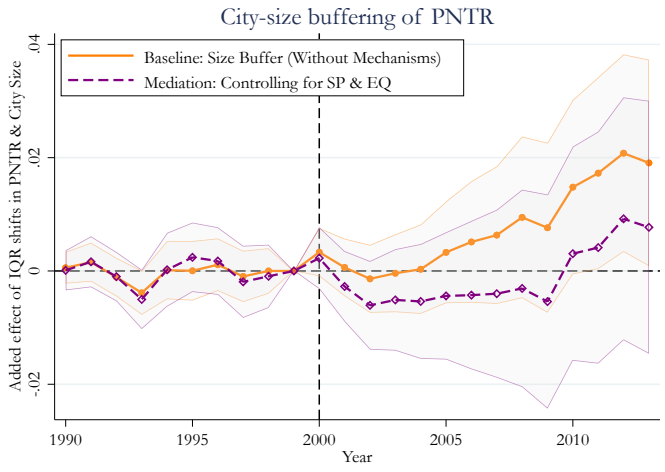


Figure 14: Mediation Analysis on the City-Size Buffering Effect. This figure plots the estimated coefficients of the triple interaction between PNTR exposure, year dummies, and city size. The solid orange line represents the baseline specification without mediators. The dashed purple line shows the remaining size effect after controlling for the interactions of PNTR with skill portability (SP) and establishment quality (EQ) as proxied by the share of city employment in establishments with above 100 employees. All coefficients are scaled to represent an IQR shift in both PNTR, city size and an respective mechanism. Shaded areas denote 95% confidence intervals. Data is from CBP, QCEW, Pierce and Schott (2020).

4.6 Falsification Tests and Robustness

As a complementary analysis, we consider two additional city characteristics—industrial diversification, proxied by the Herfindahl–Hirschman Index (HHI), and the number of firms—as alternative mediating variables in Eq. 8. The corresponding estimates are plotted in xxx. Neither measure exhibits a statistically significant moderating effect on the wage response to PNTR exposure. This finding echoes our earlier decomposition results: while diversification shapes cities' exposure to across-industry risk and firm counts help diversify idiosyncratic shocks, neither operates as an adjustment margin that mitigates the transmission of large, industry-level shocks. Taken together, the evidence points to qualitative features of local labour markets—such as skill portability and firm quality—rather than the sheer number of firms or diversification of sectors, as the key mechanisms governing cities' resilience to trade-induced disruptions.

5 CONCLUSION

This paper advances a portfolio view of cities in which urban wage growth and wage risk are jointly determined by cities' industry composition and their capacity to adjust to shocks. While initial industry weights and diversification shape cities' passive exposure to growth and risk, the analysis shows that they do not fully account for persistent cross-city differences in long-run growth-risk performance. Instead, city characteristics governing active adjustment—most notably local workers' skill portability and firm quality—play an economically meaningful role in shaping how shocks are transmitted. These findings imply that cities are not irrevocably locked into their initial portfolio positions: local adjustment capacity creates scope for improving the growth–risk profile, even for smaller cities.

The results also point to a broader research agenda. Skill portability and firm quality are unlikely to exhaust the margins of local adjustment, and other city attributes—such as financial depth, labour market institutions, housing market flexibility, or governance capacity—may also shape reallocation and risk attenuation. Beyond these empirical refinements, the urban portfolio framework developed in this paper also motivates a structural extension. Embedding the framework in a dynamic spatial general equilibrium model with heterogeneous workers and firms, mobility frictions, and location-specific adjustment technologies would allow explicit welfare analysis of cities' growth–risk trade-offs. Such a model would clarify how place-based policies reshape urban resilience and under what conditions improvements in shock absorption may come at the cost of greater inequality. Moreover, cities are not only portfolios of industries, but also bundles of housing and amenities. From an individual perspective, returns and risks associated with labour income, housing markets, and local amenities are jointly evaluated when choosing where to live. A portfolio perspective therefore might provide a unifying language for understanding location choice and place-based policies under multiple, interacting sources of risk and return.

REFERENCES

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, 109(3):659–684.

Albouy, D., Graf, W., Kellogg, R., and Wolff, H. (2016). Climate amenities, climate change, and american quality of life. *Journal of the Association of Environmental and Resource Economists*, 3(1):205–246.

Ankrim, E. M. (1992). Risk-adjusted performance attribution. *Financial Analysts Journal*, 48(2):75–82.

Anstreicher, G. (2024). Spatial influences in upward mobility. *Journal of Political Economy*, 132(1):149–199.

Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The china syndrome: Local labor market effects of import competition in the united states. *American economic review*, 103(6):2121–2168.

Bacon, C. R. (2019). *Performance Attribution: History and Progress*. CFA Institute Research Foundation.

Basak, G., Jagannathan, R., and Sun, G. (2002). A direct test for the mean variance efficiency of a portfolio. *Journal of Economic Dynamics and Control*, 26(7-8):1195–1215.

Bilal, A. and Rossi-Hansberg, E. (2021). Location as an asset. *Econometrica*, 89(5):2459–2495.

Bloom, N., Handley, K., Kurmann, A., and Luck, P. A. (2024). The china shock revisited: Job reallocation and industry switching in us labor markets. Technical report, National Bureau of Economic Research.

Brinson, G. P. and Fachler, N. (1985). Measuring non-us. equity portfolio performance. *The Journal of Portfolio Management*, 11(3):73–76.

Brinson, G. P., Hood, L. R., and Beebower, G. L. (1986). Determinants of portfolio performance. *Financial Analysts Journal*, 42(4):39–44.

Campisi, S. (2004). Debunking the interaction myth. *Journal of Performance Measurement*, 8:63–70.

Conte, M., Mejean, I., Michalski, T. K., and Schmutz, B. (2024). The volatility advantages of large labor markets. *Available at SSRN 4744092*.

Debortoli, D. and Galí, J. (2024). Idiosyncratic income risk and aggregate fluctuations. *American Economic Journal: Macroeconomics*, 16(4):279–310.

Duranton, G. and Puga, D. (2004). Micro-foundations of urban agglomeration economies. In *Handbook of regional and urban economics*, volume 4, pages 2063–2117. Elsevier.

Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica*, 79(3):733–772.

Ganong, P., Noel, P. J., Patterson, C., Vavra, J. S., and Weinberg, A. (2025). Earnings instability. Technical report, National Bureau of Economic Research.

Ganong, P. and Shoag, D. (2017). Why has regional income convergence in the us declined? *Journal of Urban Economics*, 102:76–90.

Gathmann, C. and Schönberg, U. (2010). How general is human capital? a task-based approach. *Journal of Labor Economics*, 28(1):1–49.

Gottschalk, P. and Moffitt, R. (2009). The rising instability of us earnings. *Journal of Economic Perspectives*, 23(4):3–24.

Gottschalk, P., Moffitt, R., Katz, L. F., and Dickens, W. T. (1994). The growth of earnings instability in the us labor market. *Brookings Papers on Economic Activity*, 1994(2):217–272.

Greaney, B., Parkhomenko, A., and Van Nieuwerburgh, S. (2025). Dynamic urban economics. Technical report, National Bureau of Economic Research.

Guvenen, F., Schulhofer-Wohl, S., Song, J., and Yogo, M. (2017). Worker betas: Five facts about systematic earnings risk. *American Economic Review*, 107(5):398–403.

Handley, K. and Limão, N. (2017). Policy uncertainty, trade, and welfare: Theory and evidence for china and the united states. *American economic review*, 107(9):2731–2783.

Heblich, S., Nagy, D. K., Trew, A., and Zylberberg, Y. (2025). The death and life of great british

cities. Technical report, National Bureau of Economic Research.
Hizmo, A. (2023). Risk in housing markets: An equilibrium approach. *Available at SSRN 4939847*.
Kambourov, G. and Manovskii, I. (2009). Occupational specificity of human capital. *International Economic Review*, 50(1):63–115.
Kandel, S. and Stambaugh, R. F. (1995). Portfolio inefficiency and the cross-section of expected returns. *The Journal of Finance*, 50(1):157–184.
Kikuchi, S. and O'Connor, D. G. (2025). Granular origins of agglomeration. *Available at SSRN 5454735*.
Kopytov, A., Taschereau-Dumouchel, M., and Xu, Z. (2025). The origin of risk. Working paper.
Krugman, P. (1991). *Geography and trade*. MIT press.
Levhari, D. and Weiss, Y. (1974). The effect of risk on the investment in human capital. *The American Economic Review*, 64(6):950–963.
Li, K., Sarkar, A., and Wang, Z. (2003). Diversification benefits of emerging markets subject to portfolio constraints. *Journal of empirical finance*, 10(1-2):57–80.
Manning, S. and Aguirre, T. (2025). How adaptable are american workers to ai-induced job displacement? Technical report, National Bureau of Economic Research.
Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, 7(1):77–91.
Marshall, A. (1890). 1920. principles of economics. *London: Mac-Millan*, pages 1–627.
Moretti, E. and Yi, M. (2024). Size matters: Matching externalities and the advantages of large labor markets. Technical report, National Bureau of Economic Research.
Morten, M. (2019). Temporary migration and endogenous risk sharing in village india. *Journal of Political Economy*, 127(1):1–46.
Neal, D. (1995). Industry-specific human capital: Evidence from displaced workers. *Journal of labor Economics*, 13(4):653–677.
Ortalo-Magné, F. and Prat, A. (2016). Spatial asset pricing: A first step. *Economica*, 83(329):130–171.
Overman, H. G. and Puga, D. (2010). Labor pooling as a source of agglomeration: An empirical investigation. In *Agglomeration economics*, pages 133–150. University of Chicago Press.
Parent, D. (2000). Industry-specific capital and the wage profile: Evidence from the national longitudinal survey of youth and the panel study of income dynamics. *Journal of Labor Economics*, 18(2):306–323.
Pierce, J. R. and Schott, P. K. (2016). The surprisingly swift decline of us manufacturing employment. *American economic review*, 106(7):1632–1662.
Pierce, J. R. and Schott, P. K. (2020). Trade liberalization and mortality: evidence from us counties. *American Economic Review: Insights*, 2(1):47–63.
Pierce, J. R., Schott, P. K., and Tello-Trillo, C. (2024). To find relative earnings gains after the china shock, look outside manufacturing and upstream. Technical report, National Bureau of Economic Research.
Poletaev, M. and Robinson, C. (2008). Human capital specificity: evidence from the dictionary of occupational titles and displaced worker surveys, 1984–2000. *Journal of Labor Economics*, 26(3):387–420.
Rodríguez-Clare, A., Ulate, M., and Vasquez, J. P. (2026). Trade with nominal rigidities: Understanding the unemployment and welfare effects of the china shock. *Journal of Political Economy*, 134(2):000–000.
Traiberman, S. (2019). Occupations and import competition: Evidence from denmark. *American Economic Review*, 109(12):4260–4301.
Yi, M., Mueller, S., and Stegmaier, J. (2017). Transferability of skills across sectors and heterogeneous displacement costs. *American Economic Review*, 107(5):332–336.
Yi, M., Müller, S., and Stegmaier, J. (2024). Industry mix, local labor markets, and the incidence of trade shocks. *Journal of Labor Economics*, 42(3):837–875.

CITIES AS PORTFOLIOS

ONLINE APPENDIX

Ailin Zhang*

March 13, 2026

Abstract

This online appendix contains material that complements the paper *Cities as Portfolios*. It does not replace the reading of the main paper.

*London School of Economics and Political Science. Email: a.zhang12@lse.ac.uk.

A EFFICIENT FRONTIER

This section complements Section 2.1 by providing additional details on the rationale, construction, and limitations of the urban efficient frontier. It also supplements Section 2.2 with further evidence on the efficiency gaps.

A.1 Rationale behind empirical efficient frontier

In finance, the efficient frontier is a *theoretical* object derived from a portfolio-choice problem defined on a fixed asset space: given tradable, homogeneous assets with common and fully observed return moments (μ, Σ), one can characterize the set of portfolios that maximize expected return for a given variance. In the urban context, such a structural frontier is neither empirically feasible nor well-defined. First, there is no common and fixed return–covariance structure over which cities optimize. Industry categories are not standardized assets across places: even within the same industry, cities differ in firm composition and production networks, so wage-growth processes and covariance structures are city-specific equilibrium outcomes rather than shared primitives. Moreover, the firm-level return processes and the full covariance structure among firms—objects required to construct a frontier from first principles—are not observed in the available data. Second, industrial shares are not freely chosen portfolio weights on a fixed opportunity set but equilibrium outcomes that adjust through discrete firm entry and exit, worker reallocation, and other margins involving fixed costs and thresholds. Such adjustments can reshape local technologies and shock-transmission networks, so the relevant return moments need not be invariant to changes in s; that is, (μ, Σ) cannot be treated as fixed primitives independent of the industrial structure. A Markowitz-style frontier defined by optimizing over portfolios holding (μ, Σ) constant is therefore not well-defined in our setting.

Crucially, this does *not* render the frontier concept irrelevant. Rather, it shifts its interpretation from a normative optimum to a positive benchmark: the relevant question is not "which portfolio is optimal given common primitives," but "which cities, under observed institutional and technological constraints, attain the best growth–risk combinations." The *empirical* efficient frontier therefore has substantive value as a data-driven summary of the best realized performance in the growth–risk space and as a disciplined yardstick for comparing cities' relative efficiency. Systematic dispersion inside the envelope is evidence of economically meaningful heterogeneity in inherited industrial structure and adjustment capacity, which the paper seeks to quantify and explain.

A.2 Method

We construct an *empirical* urban efficient frontier as a smoothed upper envelope of realised city outcomes in the growth–risk space, and define three gap measures that quantify a city's distance to

this envelope. Fix a period $t \in \{t_0, \ldots, t_1\}$ (e.g. 1991–2000). For each CBSA c, let $\Delta \ln w_{c,t}$ denote annual real wage growth. We summarise each city by two period-level statistics: (i) cumulative real wage growth $G_c \equiv \sum_{t=t_0}^{t_1} \Delta \ln w_{c,t}$, and (ii) realised wage risk $R_c \equiv \mathrm{sd}_t(\Delta \ln w_{c,t})$. Each city is therefore a point (R_c, G_c).

Step 1: sample upper envelope. To trace the upper envelope in the data, we sort cities by risk in ascending order and, within risk, by growth in descending order. Scanning this ordered list, we retain a city whenever its growth exceeds the running maximum among all cities with weakly lower risk. This produces a set $\mathcal{H}_1$ of cities on the *sample upper envelope*: by construction, no city with weakly lower realised risk attains strictly higher growth in the sample.

Step 2: near-envelope layer. Because an extremal envelope may contain few points and be sensitive to measurement noise, we add a second near-envelope layer. We remove $\mathcal{H}_1$ from the sample, re-apply the same envelope procedure to the remainder, and obtain $\mathcal{H}_2$. We use $\mathcal{H} \equiv \mathcal{H}_1 \cup \mathcal{H}_2$ as the set of *frontier-support* cities for smoothing (dropping duplicates when they arise).

Step 3: normalisation and spline smoothing.

Our efficiency metric is defined as a *distance* in the growth–risk plane, so the two axes must be placed on a comparable scale. We therefore rescale both variables to the unit interval within each period:

$$\tilde{R}_c = \frac{R_c - \min R}{\max R - \min R}, \qquad \tilde{G}_c = \frac{G_c - \min G}{\max G - \min G}.$$

This produces a dimensionless geometry in which distances are not mechanically driven by units and are comparable across periods.

The raw efficient frontier obtained from the frontier-support set $\mathcal{H}$ is piecewise linear and may exhibit sharp corners. To obtain a smooth approximation suitable for visualisation and distance calculations, we fit a restricted cubic spline (natural spline) $\widehat{\tilde{G}}(\tilde{R})$ using the support points $(\tilde{R}, \tilde{G}) \in \mathcal{H}$. Let $n = |\mathcal{H}|$ denote the number of frontier-support observations. To balance flexibility and stability in finite samples, we choose the number of internal knots as $K = \min\{4, \max\{2, \lfloor n/3 \rfloor\}\}$, so that the spline becomes more flexible as the frontier-support sample grows but is capped at four knots to avoid overfitting.

A restricted cubic spline approximates the frontier by representing the unknown function $\tilde{G} = f(\tilde{R})$ as a linear combination of specially constructed transformation functions of $\tilde{R}$. Given the chosen knot locations, these transformation functions, denoted $\{s_1(\tilde{R}), \ldots, s_m(\tilde{R})\}$ and often referred to as *spline basis variables*, are fully determined deterministic functions of $\tilde{R}$. They span a function space consisting of piecewise cubic functions that are twice continuously differentiable at the knots and linear beyond the boundary knots. Importantly, no estimation occurs at this stage: the basis functions are mechanically constructed once the knot locations are fixed. In practice, we implement this procedure in Stata using `mkspline s = norm_risk, cubic nknots(`K`)`, which generates the

spline basis variables from the selected knot locations (stored in `r(knots)`). We then estimate

$$\tilde{G} = \beta_0 + \sum_{\ell=1}^{m} \beta_\ell s_\ell(\tilde{R}) + \varepsilon,$$

by OLS on the frontier-support sample via `reg norm_growth s*`. Only the coefficients β_ℓ are estimated; the basis functions themselves remain fixed.

The fitted frontier is therefore

$$\widehat{\tilde{G}}(\tilde{R}) = \widehat{\beta}_0 + \sum_{\ell=1}^{m} \widehat{\beta}_\ell s_\ell(\tilde{R}).$$

For plotting and for computing distance measures, we evaluate the fitted spline on a dense grid of 400 equally spaced values of $\tilde{R}$ spanning its observed range. Using the same knot locations returned by Stata, we apply the identical spline transformation to these grid values and compute the corresponding predicted values using the estimated coefficients. This yields a smooth approximation to the empirical efficient frontier.

Step 4: gap measures. Given the estimated frontier $\widehat{\tilde{G}}(\tilde{R})$, we define three complementary notions of "distance to the frontier" for each city c.

1. *Efficiency gap (Euclidean distance).* Let $\widehat{\mathcal{F}}$ denote the piecewise-linear curve obtained from the evaluated spline grid. The efficiency gap is defined as the shortest Euclidean distance from $(\tilde{R}_c, \tilde{G}_c)$ to $\widehat{\mathcal{F}}$:

   $$\text{EffGap}_c \equiv \min_{(\tilde{r},\tilde{g})\in\widehat{\mathcal{F}}} \sqrt{(\tilde{R}_c - \tilde{r})^2 + (\tilde{G}_c - \tilde{g})^2}.$$

   Operationally, distances are computed segment by segment along the piecewise-linear frontier, allowing for orthogonal projections onto each segment; when the orthogonal projection lies outside a segment, the distance to the nearest endpoint is used.

   This measure is unit-free and captures overall inefficiency in the joint growth–risk space. It implicitly assigns equal weight to deviations in growth and risk through the Euclidean norm. Alternative norms or weighted distances could in principle reflect different relative importance of growth and risk, or be derived from an explicit welfare function. However, specifying a common welfare aggregator at the city level would require additional structural assumptions. We therefore adopt the Euclidean distance as a transparent and parsimonious benchmark.

2. *Growth gap (vertical shortfall).* The growth gap compares a city's realised growth to the frontier growth at the same risk:

   $$\text{GrowGap}_c \equiv \widehat{\tilde{G}}(\tilde{R}_c) - \tilde{G}_c.$$

   It measures how much growth a city falls short of the best observed performance *at its realised*

risk.

3. *Risk gap (horizontal excess risk).* Symmetrically, we define a risk gap as the difference between a city's realised risk and the frontier risk associated with comparable frontier growth. Operationally, we identify the frontier point whose growth is closest to $\tilde{G}_c$ and take the horizontal difference:

$$\text{RiskGap}_c \equiv \tilde{R}_c - \tilde{R}^F(\tilde{G}_c),$$

where $\tilde{R}^F(\tilde{G}_c)$ is obtained from the evaluated frontier curve.

For cities in the frontier-support set $\mathscr{H}_1$, we set all three gaps to zero by construction. Finally, for presentation we map the smoothed frontier back into the original units (R, G) and overlay it on the city scatter plots.

A.3 Limitations

The empirical frontier is a descriptive, sample-dependent object rather than a structural boundary. As an extremal envelope, it is sensitive to measurement error and may shift with additional cities or longer samples. It captures best realized performance under existing industrial structures and adjustment constraints, not a technological upper bound that cities could in principle attain. Nor is it a welfare object: it abstracts from housing costs, amenities, migration equilibria, and heterogeneous risk preferences. Throughout the paper, we therefore interpret "efficiency" as relative performance in the joint growth–risk space—distance to the empirical envelope—rather than as evidence of optimal portfolio choice or socially optimal allocations.

A.4 Additional Evidence on the Efficient Frontier

This subsection provides additional empirical evidence on the efficient frontier introduced in Section 2.1. We first document the growth–risk trade-off over the period 1990–2020, and then present the efficient frontier by decade.

Growth-risk trade-off.
While large cities tend to exhibit both higher growth and lower risk, it is natural to ask whether the urban system as a whole still displays the classic growth–risk trade-off familiar from finance. Figure A.1 shows this relationship for the pooled sample in panel (a) and by decade in panels (b)–(d). Across all CBSAs, the overall slope is positive, indicating that higher real wage growth is generally associated with greater wage volatility. One natural interpretation is that industries offering stronger growth opportunities also tend to exhibit higher intrinsic volatility in the long run, so cities more exposed to such sectors mechanically inherit both higher growth and higher

risk. In addition, the process of structural reallocation that supports faster urban growth may itself generate short- to medium-run earnings dispersion. Consistent with these forces, the relationship is economically stronger over the full 30-year horizon than within individual decades. Importantly, the decomposition in this paper shows that larger cities systematically exhibit lower within-industry wage volatility, pointing to a key channel through which they may attenuate the classic trade-off and sustain both higher growth and lower risk. Taken together, the evidence suggests that although a conventional trade-off exists in the aggregate, larger cities appear better able to attenuate—or even partially escape—it.

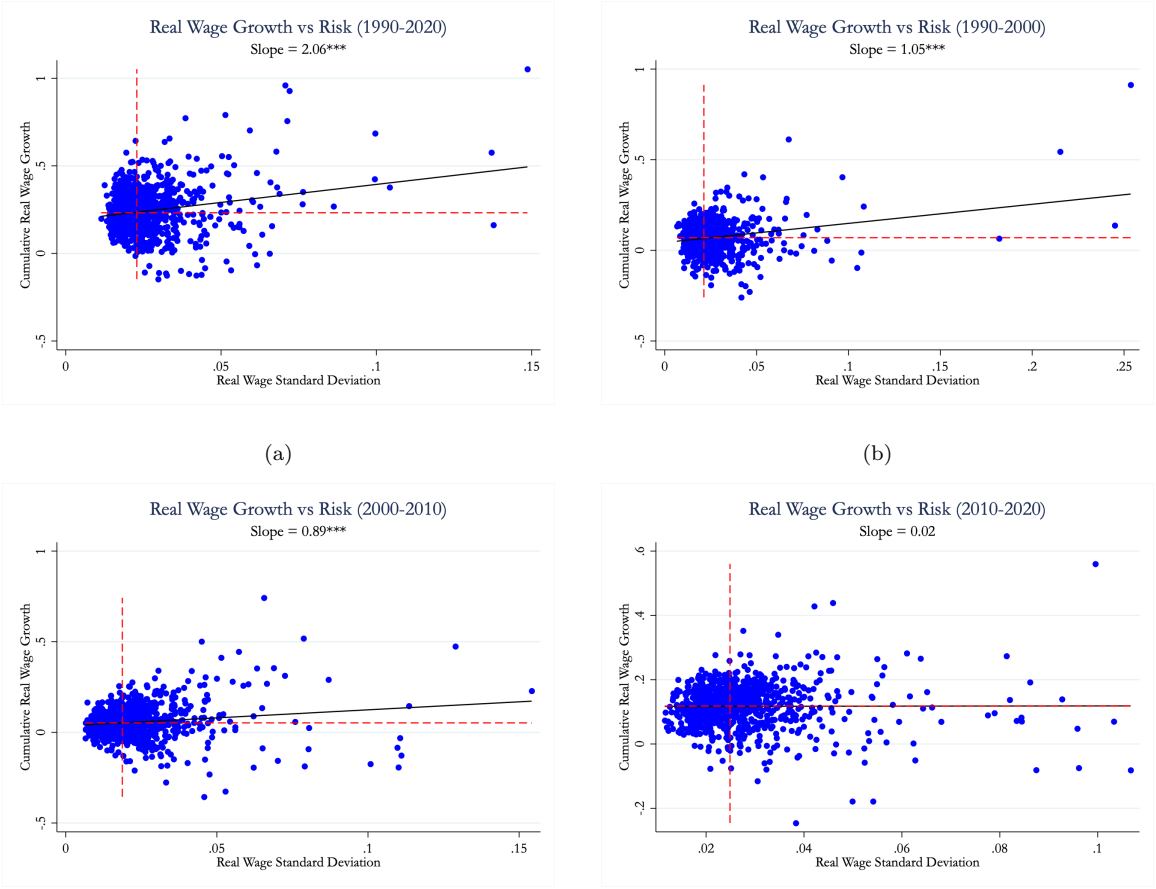


Figure A.1: Urban growth–risk trade-off across CBSAs. The top-left panel reports the full-sample frontier (1990–2020), while the remaining panels present decade-specific frontiers. The y-axis measures cumulative real wage growth over the period, and the x-axis measures the standard deviation of real wage growth over the corresponding time series. The data are private-sector average annual pay from the QCEW, deflated using the CPI-U.

Efficient frontier and density elasticity of efficiency by decade.

While Figures 1 and 2 summarize the efficient frontier and the density elasticity of efficiency gaps over 1990–2020, Figure A.2 reports the same evidence by decade for all CBSAs. Panels (a)–(c) highlight substantial time variation in the performance of individual cities: even cities that are

highly efficient in the long run—such as San Francisco, Seattle, and San Jose—do not consistently lie near the frontier in every decade. This instability is especially visible in the 2000s, a period marked by large aggregate and sectoral shocks (e.g., the post–dot-com adjustment early in the decade and the Great Recession at the end), which likely amplified wage-growth volatility in tech-intensive metros. For instance, San Jose exhibits notably high wage-growth volatility in the 2000s. Despite these fluctuations in city rankings over time, the cross-sectional pattern by size is remarkably stable. Panel (d) shows a clear density gradient in efficiency gaps: the coefficients from regressing each efficiency-gap measure on log population are significantly negative in all three decades. This implies that, within each period, larger CBSAs are systematically closer to the efficient frontier and therefore more efficient in a joint growth–risk sense.

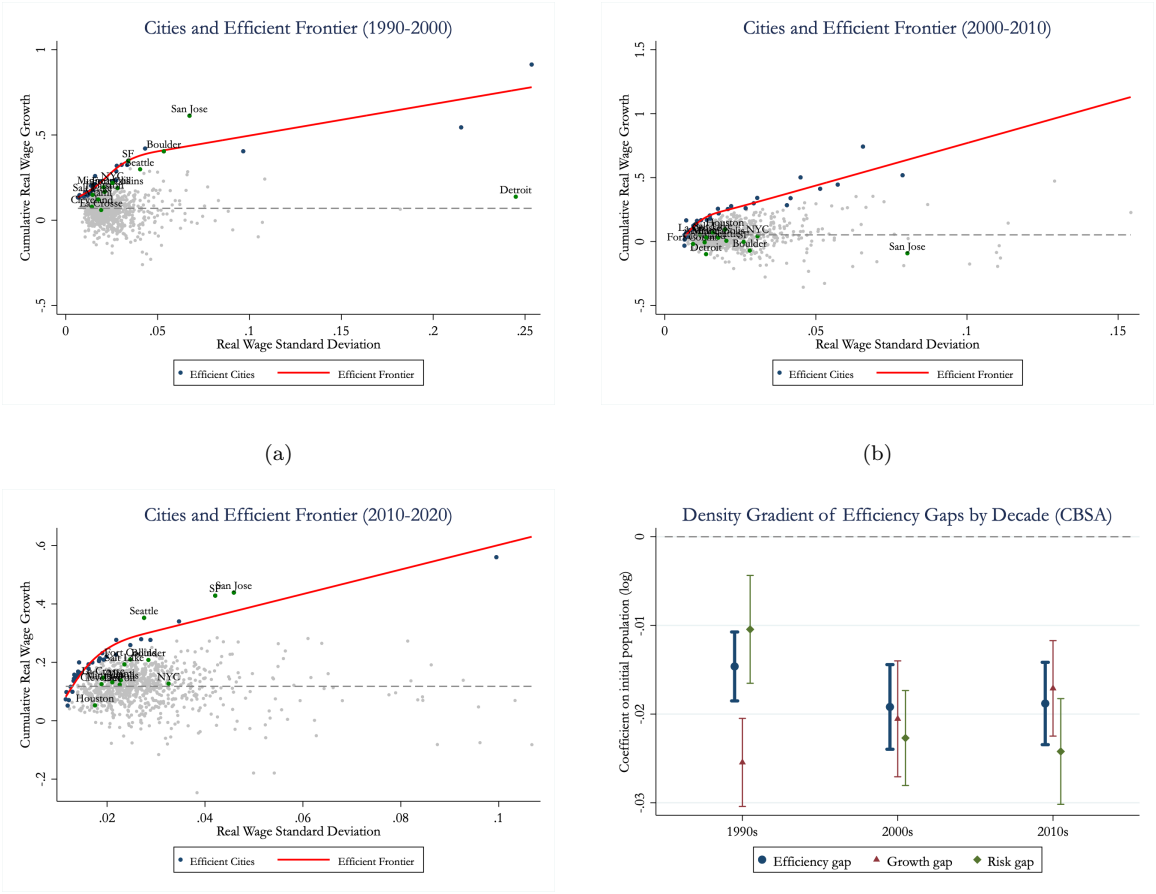


Figure A.2: Urban growth–risk efficient frontiers and associated efficiency gap gradients by decade. Panels (a)-(c) report decade-specific frontiers (1990–2000, 2001–2010, and 2011–2020). Panel (d) shows the corresponding density gradients of the efficiency gaps estimated using Eq. 2.2. Data are private-sector average annual pay from the QCEW, deflated using the CPI-U.

B DECOMPOSITION

B.1 Variance Approximation

This section clarifies (i) the object that (2.3.1) decomposes exactly and (ii) the conditions under which that object is a good proxy for the time-series volatility of city-average wage growth.

Let $G_{c,t} \equiv s'_{c,t} g_{c,t}$ denote city c's employment-share-weighted real wage growth in year t, where $s_{c,t} \in \Delta^{J-1}$ is the vector of realised industry employment shares and $g_{c,t} \in \mathbb{R}^J$ the vector of industry wage-growth rates. Over a horizon T, let $\mathbf{\Omega}_T^N$ and $\mathbf{\Omega}_{c,T}$ denote the national and city-specific variance-covariance matrices of $g_{c,t}$. Define the quadratic-form object

$$\mathcal{R}_{c,T} \;\equiv\; \overline{s'_{c,t} \mathbf{\Omega}_{c,T} s_{c,t}}. \tag{B.1}$$

The decomposition in (2.3.1) is *algebraically exact* for $\mathcal{R}_{c,T}$: using $\mathbf{\Omega}_{c,T} = \mathbf{\Omega}_T^N + (\mathbf{\Omega}_{c,T} - \mathbf{\Omega}_T^N)$ and adding–subtracting $s'_{c,0} \mathbf{\Omega}_T^N s_{c,0}$,

$$\mathcal{R}_{c,T} = \underbrace{s'_{c,0} \mathbf{\Omega}_T^N s_{c,0}}_{\text{Benchmark risk exposure}} + \underbrace{\left(\overline{s'_{c,t} \mathbf{\Omega}_T^N s_{c,t}} - s'_{c,0} \mathbf{\Omega}_T^N s_{c,0} \right)}_{\text{Reallocation risk}} + \underbrace{\overline{s'_{c,t} (\mathbf{\Omega}_{c,T} - \mathbf{\Omega}_T^N) s_{c,t}}}_{\text{Idiosyncratic risk}}. \tag{B.2}$$

Accordingly, the "$\approx$" in (2.3.1) comes from interpreting $\mathcal{R}_{c,T}$ as an approximation for the time-series volatility of $G_{c,t}$.

Interpreting $\mathcal{R}_{c,T}$ as an approximation of $\mathrm{Var}_T(G_{c,t})$ rests on a few assumptions. Write industry wage growth as the sum of a within-window mean and a zero-mean innovation,

$$g_{c,t} = \bar{g}_{c,T} + u_{c,t}, \qquad \overline{u_{c,t}} = \mathbf{0}, \qquad \overline{u_{c,t} u'_{c,t}} = \mathbf{\Omega}_{c,T}, \tag{B.3}$$

Then

$$G_{c,t} = s'_{c,t} \bar{g}_{c,T} + s'_{c,t} u_{c,t}. \tag{B.4}$$

Under the orthogonality condition that innovations have zero conditional mean given contemporaneous shares, $\mathbb{E}[u_{c,t} \mid s_{c,t}] = \mathbf{0}$, we have $\mathbb{E}[s'_{c,t} u_{c,t}] = 0$, so that $\overline{s'_{c,t} u_{c,t}}$ is approximately zero in large samples. Using $\mathrm{Var}_T(G) = \overline{(G - \bar{G})^2}$ and the decomposition above, the gap can be written as

$$\mathrm{Var}_T(G_{c,t}) - \mathcal{R}_{c,T} = \left(\mathrm{Var}_T(s'_{c,t} u_{c,t}) - \overline{s'_{c,t} \mathbf{\Omega}_{c,T} s_{c,t}} \right) + 2\,\mathrm{Cov}_T\!\left(s'_{c,t} \bar{g}_{c,T},\, s'_{c,t} u_{c,t}\right) + \mathrm{Var}_T(s'_{c,t} \bar{g}_{c,T}), \tag{B.5}$$

where the three terms on the right-hand side capture, respectively, (i) the quadratic-form approximation error for innovations, (ii) share–innovation comovement within the window, and (iii) variation induced by changing exposure to within-window mean growth rates.

A sufficient set of conditions for $\mathrm{Var}_T(G_{c,t}) \approx \mathcal{R}_{c,T}$ is therefore:

(C1) *Second-moment stability of innovations.* Let $\mathbf{\Omega}_{c,t} \equiv \mathbb{E}[u_{c,t}u'_{c,t} \mid s_{c,t}]$ denote the period-specific innovation covariance matrix. Because $\mathbf{\Omega}_{c,t}$ is not directly observable in annual data, the empirical implementation relies on a window-level estimate $\mathbf{\Omega}_{c,T}$ constructed from realised innovations. Condition (C1) requires that second moments are approximately stable within the window so that $\mathbf{\Omega}_{c,t} \approx \mathbf{\Omega}_{c,T}$. Under weak dependence, this implies that $\mathrm{Var}_T(s'_{c,t}u_{c,t})$ is well approximated by $\overline{s'_{c,t}\mathbf{\Omega}_{c,T}s_{c,t}}$.

(C2) *Limited share–shock endogeneity.* Innovations are orthogonal to contemporaneous shares in the sense that $\mathbb{E}[u_{c,t} \mid s_{c,t}] = \mathbf{0}$. Since $s'_{c,t}\bar{g}_{c,T}$ is measurable with respect to $s_{c,t}$, this implies

$$\mathrm{Cov}_T\big(s'_{c,t}\bar{g}_{c,T},\, s'_{c,t}u_{c,t}\big) \approx 0,$$

so term (2) is negligible. Economically, this condition requires that unexpected wage-growth innovations are not systematically predictable from the contemporaneous industry composition of the city.

(C3) *Limited mean-growth exposure variation.* Within window T, $\bar{g}_{c,T}$ is time-invariant, hence

$$\mathrm{Var}_T(s'_{c,t}\bar{g}_{c,T}) = \bar{g}'_{c,T}\,\mathrm{Var}_T(s_{c,t})\,\bar{g}_{c,T} \;\le\; \|\bar{g}_{c,T}\|^2 \cdot \lambda_{\max}\big(\mathrm{Var}_T(s_{c,t})\big),$$

where $\lambda_{\max}(\cdot)$ denotes the largest eigenvalue of the share variance-covariance matrix, i.e. the maximal within-window variation of employment shares along any direction. This expression shows that this term is governed by two objects: (i) dispersion in within-window mean industry growth rates $\bar{g}_{c,T}$, and (ii) the magnitude of within-window share fluctuations. Economically, this term captures variation generated when employment shares gradually rotate toward industries with persistently higher or lower mean growth rates. This is conceptually distinct from the risk object in $\mathcal{R}_{c,T}$, which captures exposure to high-frequency wage-growth innovations through the covariance structure $\Omega_{c,T}$. In annual data, industry shocks typically generate sizeable year-to-year wage fluctuations, whereas employment shares evolve more slowly due to adjustment frictions, entry and exit dynamics, and labour reallocation costs. Consequently, changes in $s'_{c,t}\bar{g}_{c,T}$ represent a low-frequency component of variation relative to innovation-driven volatility. When this low-frequency component is small relative to $\mathcal{R}_{c,T}$, the quadratic-form object provides a close proxy for innovation-driven city-level wage risk.

Importantly, the three terms on the right-hand side of (B.5) need not be exactly zero. Conditions (C1)–(C3) instead require that each term is of second-order importance relative to $\mathrm{Var}_T(G_{c,t})$, so their joint contribution is negligible. Under these conditions, the quadratic-form object $\mathcal{R}_{c,T}$ provides a close proxy for the volatility of realised city-average wage growth over window T.

B.2 Idiosyncratic Variance and Covariance

This section complements the analysis of idiosyncratic risk in Section 2.3 of the main paper. Section 2.3.1 shows that the idiosyncratic risk component—driven by evolving industry weights and city-specific deviations in the variance–covariance structure from the national benchmark in Eq. 2.3.1—is the primary source of lower wage risk in larger cities. Section 2.3.2 further demonstrates that this advantage is mainly attributable to a lower aggregate contribution of within-industry variance rather than reduced between-industry covariance, as illustrated in Figure 5. Nevertheless, cities differ in their constituent industries, raising two related questions: for a given industry, is wage growth more volatile in large cities; and, more importantly, for a given industry pair, is their covariance higher or lower in large cities? While the former seems straightforward, the latter remains largely unexplored in the existing literature.

Though existing work emphasizes how production networks generate aggregate co-movement across industries (Dupor, 1999; Shea, 2002; Conley and Dupor, 2003; Acemoglu et al., 2012; Carvalho, 2014), it largely abstracts from spatial heterogeneity. As a result, whether the covariance of a given industry pair is systematically lower or higher in larger cities remains theoretically ambiguous. On the one hand, larger cities may feature more localized input–output linkages, with upstream and downstream suppliers more likely to be colocated, which can amplify the propagation of sectoral shocks and increase cross-industry co-movement. On the other hand, larger cities also offer thicker markets and greater adjustment capacity—through diversified supplier and customer bases, within- and across-firm reallocation, and entry and exit on the extensive margin—which can dampen shock transmission and weaken pairwise covariance even when technological linkages are strong.

Empirically, to address the first question—whether a given industry is less volatile in larger cities—we estimate the following city–industry–decade level regression:

$$\mathrm{Var}_{ic} = \beta \, \ln L_c \; + \; \mu_i \; + \; \varepsilon_{ic},$$

where Var_{ic} denotes the within-industry variance of real wage growth for industry i in city c, $\ln L_c$ is log initial city employment, and μ_i captures industry fixed effects. Across all three decades, the estimated coefficient on city size is negative and highly statistically significant, indicating that for a given industry, wage growth is systematically less volatile in larger cities. This result provides complementary evidence that large cities mitigate idiosyncratic risk primarily through lower within-industry volatility discussed in Section 2.3.2.

Then more importantly, motivated by the theoretical ambiguity related to covariance, we examine how industry-level co-movement varies with city size by estimating a city–industry-pair regression of the form:

$$\mathrm{Cov}_{ijc} = \beta \ln L_c + \mu_{ij} + \varepsilon_{ijc},$$

where Cov_{ijc} denotes the covariance of real wage growth between industries i and j in city c, and $\ln L_c$

Table B.1: Within-industry wage variance and city size, by decade

	1990s	2000s	2010s
Log employment	-0.00387^{***}	-0.00463^{***}	-0.00320^{***}
	(0.00020)	(0.00029)	(0.00050)
Observations	71,510	107,446	107,896
R^2	0.074	0.140	0.075

Notes: The dependent variable is the within-industry variance of real wage growth. Each observation corresponds to an industry–city–decade cell. All regressions include industry fixed effects. Standard errors are clustered at the CBSA level.
Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

is initial city employment. Industry-pair fixed effects μ_{ij} absorb pair-specific national co-movement, so identification comes from within-pair variation across cities of different sizes. Standard errors are clustered at the city level.

Panel (A) of Table B.2 reports estimates under unweighted regression (coefficients are scaled by 10^5 for visualization). City size is positively associated with industry-pair covariance in the 1990s and 2010s, but not in the 2000s, implying no stable pattern of systematically higher or lower co-movement for a given pair in larger cities. A plausible interpretation is that large cities feature denser input–output linkages that can amplify co-movement, yet this channel is muted in the 2000s when the covariance structure was reshaped by broad, common shocks—notably the China shock and the housing-finance cycle—whose incidence was both more national in scope and more heterogeneous across cities in ways not tightly aligned with initial size. Panel (B) weights pairs by their initial employment-share product, focusing on economically important pairs; under this specification the size–covariance relationship further attenuates and becomes statistically insignificant. Overall, these results suggest that large cities' lower aggregate wage risk is unlikely to operate through pervasive dampening of pairwise covariance, and instead aligns with the decomposition evidence that the risk advantage primarily reflects lower within-industry volatility and compositional forces rather than systematic attenuation of cross-industry co-movement at the pair level.

Table B.2: Industry-pair covariance and city size, by decade

(A) Unweighted

	1990s	2000s	2010s
Log employment	2.270***	-0.050	1.880*
	(0.443)	(0.677)	(0.971)
Observations	4,136,778	7,425,279	7,515,291
R^2	0.014	0.021	0.014

Notes: Dependent variable is the industry-pair covariance of real wage growth in each decade, scaled by 10^5 for visualization. All specifications include industry-pair fixed effects (industry × industry). Standard errors are clustered at the CBSA level. *Significance:* * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

(B) Weighted by pair importance

	1990s	2000s	2010s
Log employment	1.690*	-0.381	0.758
	(0.916)	(0.862)	(0.760)
Observations	4,136,778	7,425,279	7,515,291
R^2	0.038	0.028	0.047

Notes: Dependent variable is the industry-pair covariance of real wage growth in each decade, scaled by 10^5 for visualization. Analytical weights are defined as the product of the initial employment shares of each industry pair, $w_{ij} = \bar{s}_{ic} \times \bar{s}_{jc}$. All specifications include industry-pair fixed effects (industry × industry). Standard errors are clustered at the CBSA level. *Significance:* * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

B.3 A Heuristic Framework

This section introduces a stylized, partial-equilibrium heuristic framework to formalize the intuition behind our empirical mechanism tests in Section 4 of the main paper. Rather than attempting to structurally model the complex, long-run transitional dynamics of urban structural transformation, this toy model is specifically designed to isolate how local frictions process a given external shock. By abstracting to a representative shock-and-adjustment cycle, we aim to disentangle a city's static *baseline exposure*—anchored by inherited diversification and firm counts—from its active *dynamic shock-absorption capacity*, which is governed by two distinct local adjustment margins: (1) within-industry buffering via firm organizational capacity, and (2) short-run labour reallocation via worker skill portability.

B.3.1 Baseline Exposure

Consider a representative cycle where city c has an inherited industry employment share vector $s_{c,0}$. Each industry i faces a national structural shock $\varepsilon_i \sim (0, \sigma_i^2)$ drawn from a stationary national covariance matrix Σ. Concurrently, each local firm f faces an independent, idiosyncratic disturbance $\eta_{fic} \sim (0, \sigma_\eta^2)$.

The city's expected baseline growth is $\bar{g}_c^{base} = s'_{c,0}\mu$, where μ denotes deterministic national trends. Before any active adjustment takes place, the city's *passive* risk profile is dictated by its initial composition. The exposure to systematic industry shocks generates a baseline covariance risk of $s'_{c,0}\Sigma s_{c,0}$. Meanwhile, exposure to idiosyncratic firm-level noise (η_{fic}) is mechanically smoothed by the city's granular production structure. If $h_{ic} \equiv \sum_{f=1}^{M_{ic}} w_{fic}^2$ is the Herfindahl-Hirschman Index (HHI) of local firm concentration, the variance of the pooled idiosyncratic component is $h_{ic}\sigma_\eta^2$. Importantly, a highly fragmented firm base (a large number of firms M_{ic}, driving $h_{ic} \to 0$) mechanically

diversifies away idiosyncratic firm-level noise. However, this theoretical distinction clarifies why mechanical diversification—whether across industries or across firms within an industry—shapes static baseline exposure but does not act as a buffer against structural shocks (ε_i).

B.3.2 Stage 1: Within-industry Firm Buffering

When the national industry-wide disturbance ε_i is transmitted to city c through local industry i, its local pass-through is actively mitigated by the local industry's firm-side buffering capacity, summarized by $\theta_{ic} \in [0, 1)$.

To micro-found this parameter, consider a local firm f with operational scale S_f. Motivated by Kopytov et al. (2025), which emphasizes the scale advantage in corporate risk management, we allow the firm to actively choose its buffering capacity θ_f to minimize the joint burden of shock exposure and risk mitigation. Unbuffered volatility imposes a disruption loss that scales proportionally with firm size: $\frac{1}{2}\lambda S_f(1-\theta_f)^2\sigma_i^2$, where $\lambda > 0$. Conversely, maintaining buffering capacity entails an organizational resource cost: $\frac{1}{2}(F + vS_f)\theta_f^2$. Here, $v > 0$ captures variable costs, while $F > 0$ represents the fixed organizational overhead. The firm solves:

$$\min_{\theta_f} \left\{ \frac{1}{2}\lambda S_f(1-\theta_f)^2\sigma_i^2 + \frac{1}{2}(F + vS_f)\theta_f^2 \right\} \tag{B.6}$$

The optimal interior solution for the firm's buffering capacity is:

$$\theta_f^* = \frac{1}{1 + \frac{v}{\lambda\sigma_i^2} + \frac{F}{\lambda S_f \sigma_i^2}}, \tag{B.7}$$

which analytically formalizes the scale advantage in risk management. Because of the positive fixed organizational cost ($F > 0$), larger firms benefit from economies of scale. Consequently, the optimal buffering capacity θ_f^* is strictly increasing in firm scale S_f.

By aggregating individual responses, the effective buffering capacity at the city-industry level emerges as the employment-weighted average of firm-level capacities:

$$\theta_{ic} \equiv \sum_{f=1}^{M_{ic}} w_{fic}\theta_f^*(S_f) \tag{B.8}$$

In this setup, θ_{ic} formalizes the capacity to buffer structural shocks, which is driven by operational scale. Empirically, we interpret this parameter as capturing broader organizational and financial resilience. Consequently, the pre-reallocation latent wage growth is modelled as:

$$g_{ic}^{pre} = \mu_i + (1-\theta_{ic})\varepsilon_i + \bar{\eta}_{ic}, \tag{B.9}$$

capturing how the pass-through of structural industry shocks is shaped by the local firm base's

buffering capability (θ_{ic}).

Proposition 1 (Firm Scale and Aggregate Buffering Capacity) *The aggregate shock-buffering capacity of a local industry (θ_{ic}) increases in the employment-weighted buffering capacity of its constituent firms. In this setup, that buffering capacity is itself increasing in firm scale. Consequently, while the sheer fragmentation of firms (low h_{ic}) mechanically smooths idiosyncratic noise, the prevalence of large, high-capacity establishments systematically endows the local industry with a stronger buffer against national industry-wide shocks.*[1]

B.3.3 Stage 2: Reallocation, Loss Dispersion, and Local Demand Multipliers

Following Stage 1, let the pre-reallocation wage pressure in local industry i be

$$g_{ic}^{pre} = \mu_i + (1 - \theta_{ic})\varepsilon_i, \tag{B.10}$$

where an adverse structural shock implies a strongly negative g_{ic}^{pre} in exposed sectors. Displaced workers attempt to mitigate this wage loss by reallocating to other local sectors. However, their mobility is bounded by occupation-specific human capital. Let the fraction of workers successfully exiting the shocked industry i be

$$\xi_{ic} = \tau_i \, m_{ic} \, (-g_{ic}^{pre})_+, \qquad m_{ic} \equiv \sum_k \phi_{ikc}(1 - s_{ic|k}), \tag{B.11}$$

where m_{ic} captures the thickness of feasible within-occupation local destinations, ϕ_{ikc} is the skill composition of industry i, and $\tau_i > 0$ governs the responsiveness of labour outflows to adverse wage pressure. The corresponding outflow, expressed as a share of total city employment, is $x_{ic} = s_{ic}\xi_{ic}$.

Let A_{jic} denote the share of outflows from origin industry i absorbed by destination industry j, with $\sum_{j \neq i} A_{jic} = 1$. The net labour-supply change in sector j is

$$n_{jc} = \sum_i A_{jic} x_{ic} - x_{jc}. \tag{B.12}$$

Assuming downward-sloping local labour demand, post-reallocation wage growth satisfies

$$g_{jc}^{post} = g_{jc}^{pre} - \kappa_j \frac{n_{jc}}{s_{jc}}, \tag{B.13}$$

where $\kappa_j > 0$ is the inverse labour-demand elasticity.

[1] In smaller local markets, limited total scale may sharpen the trade-off between buffering structural shocks through larger establishments and smoothing idiosyncratic noise through firm fragmentation; larger cities are more likely to relax this constraint.

Redistributing Wage Losses. Holding initial employment weights fixed, city-wide average wage growth is approximately

$$\bar{g}_c^{post} \approx \bar{g}_c^{pre} + \sum_i x_{ic} \left[\kappa_i - \sum_j \kappa_j A_{jic} \right], \tag{B.14}$$

where $\bar{g}_c^{pre} = \sum_j s_{jc} g_{jc}^{pre}$. While the direct price effect on the arithmetic average wage depends on relative sectoral elasticities and is therefore theoretically ambiguous, Equation (4) reveals a more robust mechanism: skill portability reduces the concentration of adverse wage pressure. Outflows attenuate the deepest wage declines in directly hit sectors, at the cost of milder wage compression spread across a broader set of receiving sectors.

Weakening the Local Demand Multiplier. This redistribution becomes aggregate-wage stabilizing once local spillovers are admitted. Empirical literature shows that concentrated structural shocks propagate through local economies via demand multipliers, depressing wages even in unexposed non-tradable sectors Autor et al. (2013). Suppose this local secondary damage is increasing and convex in the concentration of sectoral income losses, as extreme localized wage drops trigger non-linear contractions such as mass defaults and service-sector bankruptcies. Then dispersing a given adverse shock more broadly across sectors reduces aggregate multiplier damage relative to allowing a deep collapse in a narrow set of industries. In this sense, skill portability does not primarily operate by raising first-order average wage growth directly; rather, it stabilizes city-wide wages by preventing concentrated sectoral contractions from propagating through local demand.

Proposition 2 (Skill Portability and Aggregate Resilience) *Following an adverse sectoral demand shock, higher local skill portability increases the scope for within-occupation reallocation across sectors. Although the direct price effect on the city-wide average wage may be ambiguous, reallocation reduces the concentration of wage losses in the most exposed local sectors. When local secondary contractions are increasing and convex in the concentration of sectoral income losses, higher skill portability stabilizes aggregate city-wide wages by weakening the local demand multiplier.*

C DATA

This section complements Section 3 in the main paper by providing more details on the data used and variables constructed.

C.1 Wage and Employment (QCEW)

Private-sector focus. The baseline analysis restricts attention to private-sector employment and wages in QCEW. Our objects of interest—industry exposure, reallocation, and covariance structure—are intended to capture market-driven wage risk and structural adjustment to industry shocks.

Public-sector employment is governed by institutional and fiscal rules, and its wage and employment dynamics are not directly comparable to those in competitive industries. Including government would therefore mix distinct wage processes and mechanically dampen measured volatility in cities with large public sectors. Focusing on the private sector yields a consistent industry mapping and a cleaner portfolio interpretation of city outcomes. As a robustness check, the Supplement reports results including both private and public sectors; the main conclusions remain unchanged.

Employment share. QCEW industry employment data are subject to nondisclosure rules imposed by the Bureau of Labor Statistics (BLS). At the county–industry level, cells are suppressed when disclosure criteria are not met, for example if fewer than three establishments are present or if a single establishment accounts for more than 80 percent of employment. Although detailed cells may be suppressed, BLS publishes control totals at higher levels of aggregation that incorporate the suppressed employment. Consequently, observed industry employment at finer NAICS levels generally does not sum to total private employment, particularly in smaller counties. Let total private employment in city c and year t be denoted by E_{ct}^{tot}, and observed industry employment be E_{cit}^{d} for industry $i \in \mathcal{I}_{ct}^{obs}$ at digit level d. Typically,

$$\sum_{i \in \mathcal{I}_{ct}^{obs}} E_{cit}^{d} \;\leq\; E_{ct}^{tot}. \tag{C.15}$$

To conduct the decomposition analysis, consider two possible constructions of industry shares:

(i) Renormalized shares (discarding missing mass).

$$\tilde{s}_{cit}^{d} = \frac{E_{cit}^{d}}{\sum_{j \in \mathcal{I}_{ct}^{obs}} E_{cjt}^{d}}, \qquad \sum_{i \in \mathcal{I}_{ct}^{obs}} \tilde{s}_{cit}^{d} = 1. \tag{C.16}$$

(ii) Shares relative to total employment.

$$s_{cit}^{d} = \frac{E_{cit}^{d}}{E_{ct}^{tot}}, \qquad \sum_{i \in \mathcal{I}_{ct}^{obs}} s_{cit}^{d} \leq 1. \tag{C.17}$$

We adopt approach (ii). Renormalization mechanically rescales observed industry weights by the inverse of the observed employment share. Specifically, the renormalized weights is

$$\tilde{s}_{cit}^{d} = \frac{s_{cit}^{d}}{\sum_{j \in \mathcal{I}_{ct}^{obs}} s_{cjt}^{d}}.$$

Consequently, quadratic portfolio measures such as $\tilde{s}_{ct}^{d\prime} \Omega \tilde{s}_{ct}^{d}$ are mechanically inflated relative to $s_{ct}^{d\prime} \Omega s_{ct}^{d}$ by a factor of $\left(\sum_{j \in \mathcal{I}_{ct}^{obs}} s_{cjt}^{d} \right)^{-2}$. Because the observed employment share differs systematically across cities and varies over time, renormalization might mechanically inflate measured risk in

locations with more suppression and induces artificial changes in industry weights when the set of observed industries changes. This generates both cross-sectional distortion and spurious reallocation in decomposition exercises that rely on variation in shares.

More fundamentally, industry weights in a portfolio framework must be defined relative to total employment. Renormalization implicitly reallocates the employment mass of suppressed industries across observed sectors, thereby altering the economic object under study. Defining shares relative to total private employment preserves the true scale of the city economy and treats suppressed industries as an unobserved residual component.

We therefore construct industry shares relative to total employment and do not impute suppressed cells. Because suppression is more prevalent in smaller cities—particularly at finer industry levels—the portfolio-based risk measure $\hat{R}^d_{ct} = s^{d\prime}_{ct}\Omega s^d_{ct}$ may understate total wage volatility in cities with a larger fraction of suppressed employment. To mitigate this attenuation in cross-city comparisons, we rescale the estimated risk components so that their sum matches the directly observed variance of city-level real wage growth. Let $\mathrm{Var}(\Delta w_{ct})$ denote the observed city wage variance, and $\sum_k \hat{R}^d_{ct,k}$ the sum of estimated risk components. We define the adjusted components as

$$\hat{R}^{adj}_{ct,k} = \hat{R}^d_{ct,k} \cdot \frac{\mathrm{Var}(\Delta w_{ct})}{\sum_k \hat{R}^d_{ct,k}}.$$

This preserves the relative contribution of each component while aligning total portfolio risk with observed wage volatility. In addition, we report results at the 3-digit and 2-digit NAICS levels alongside the 4-digit baseline. Because suppression is less severe at higher levels of aggregation, these estimates provide robustness to potential attenuation from missing cells.

C.2 Compustat.

This subsection complements Section 3.3 by providing additional information on the variables complied from Compustat dataset.

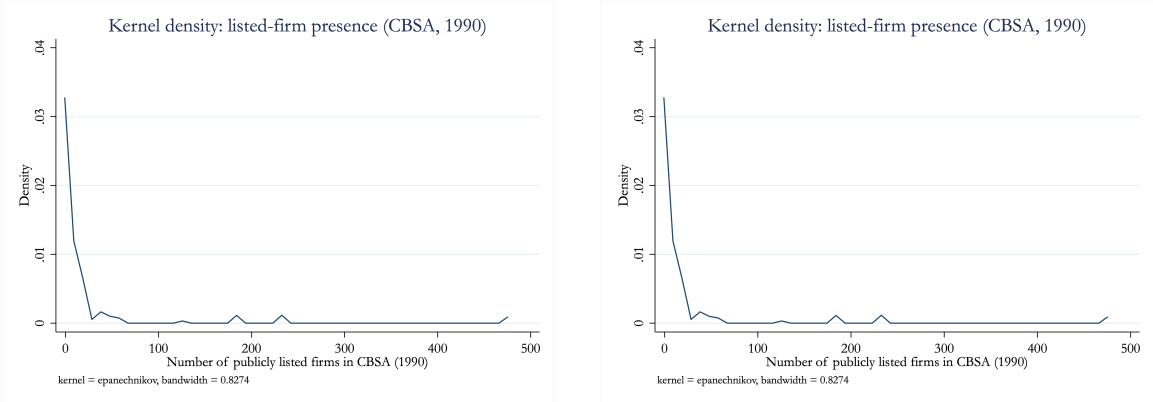
Compustat is a restricted database accessed via Wharton Research Data Services (WRDS), providing comprehensive financial and market information on publicly listed firms worldwide. It covers approximately 99% of global market capitalization and includes annual firm-level data dating back to 1950. We draw on the Compustat Fundamental Annual dataset and the Historical Segments files for North America to construct a range of firm-level quality measures, which we subsequently aggregate to the CBSA level. Firms in Fundamental Annual database but not in the Historical Segment database are treated as single-business-segment firms. Table C.3 details the construction of the variables used in this study.

We restrict the sample to firms with stock ownership code `STKO`=0, corresponding to companies listed on major exchanges (NYSE, NASDAQ, etc.), in order to focus on established, high-quality

publicly traded firms that are more likely to represent the leading corporate presence in a city. Firms traded over-the-counter (STKO=3), classified as subsidiaries (STKO=1/2), or undergoing leveraged buyouts (STKO=4) are typically smaller, less economically central, or characterized by atypical capital structures, and may not capture the type of corporate strength relevant for shock buffering. We map ZIP codes to Core-Based Statistical Areas (CBSAs) using the HUD–USPS ZIP Code Crosswalk published by the U.S. Department of Housing and Urban Development's Office of Policy Development and Research (PD&R).[2]. Around xxx% of publicly listed firms are located in MSAs, while xxx% are located in μSAs, consistent with Davis and Henderson (2008) that MSAs are usually clusters of firm headquarters.[3] Figure C.3 displays the distribution of listed firms across CBSAs.

Although Compustat provides standardized financial and market information, its use in our analysis is subject to several limitations. First, the dataset primarily covers publicly traded firms and thus reflects the right tail of the firm-size distribution rather than average firm quality. Second, firm location is identified using the ZIP code of corporate headquarters, which may not fully capture the geographic distribution of firm activities, particularly for multi-establishment firms. Third, the employment measure reflects total firm-level employment rather than employment located in the headquarters CBSA. Although this may introduce measurement error for multi-establishment firms, total firm employment is a salient proxy for firm scale and economic prominence, so employment-weighted city aggregation still captures the characteristics of the leading corporate sector anchored in the city.

Figure C.3: Listed Firms and City Size (1990)



(a) Kernel density of listed-firm count

(b) Listed firms and city employment

Notes: Panel (a) shows the kernel density of the number of Compustat-listed firms across CBSAs in 1990. Panel (b) plots the number of listed firms against total city employment in 1990. Listed firms are assigned to CBSAs based on zip code of companies' headquarters. City employment is measured using total private-sector employment with QCEW data.

[2] `https://www.huduser.gov/portal/datasets/usps_crosswalk.html`

[3] Metropolitan Statistical Areas (MSAs) and Micropolitan Statistical Areas (μSAs) are statistical regions defined by the Office of Management and Budget; MSAs contain an urban core with a population of at least 50,000, whereas μSAs are centered on smaller urban cores with populations between 10,000 and 49,999.

Table C.3: Variables Constructed from Compustat (Fiscal Year 1990)

Variable	Definition	Abbreviations	Economic Dimension
Panel A. From Compustat Fundamentals Annual database			
Listed-firm employment share	$\frac{\sum_{f \in c} EMP_f}{EmpTotal_c}$	Share of city employment accounted for by publicly traded firms.	Corporate intensity
Leverage	$\frac{DLC + DLTT}{AT}$	DLC: current debt; DLTT: long-term debt; AT: total assets.	Financial resilience
Interest coverage	$\frac{EBITDA}{XINT}$	EBITDA: earnings before interest, taxes, depreciation, amortization; XINT: interest expense.	Financial resilience
Cash-to-assets	$\frac{CHE}{AT}$	CHE: cash and short-term investments.	
Return on assets	$\frac{IB}{AT}$	IB: income before extraordinary items.	Profitability & efficiency
Operating margin	$\frac{OIADP}{SALE}$	OIADP: operating income after depreciation; SALE: net sales.	Profitability & efficiency
Asset turnover	$\frac{SALE}{AT}$		Profitability & efficiency
Capital expenditure intensity	$\frac{CAPX}{AT}$	CAPX: capital expenditures.	Investment capacity
R&D intensity	$\frac{XRD}{SALE}$	XRD: research and development.	Investment capacity
PPE-to-assets	$\frac{PPENT}{AT}$	PPENT: property, plant, and equipment.	Asset structure
Panel B. From Compustat Historical Segments files			
Business-segment count	# segments	Business segments reported in Compustat Historical Segments.	Business diversification
Industry breadth	# distinct SIC codes		Business diversification
Segment sales Herfindahl index	$\sum_s \left(\frac{sales_s}{\sum_k sales_k} \right)^2$	Based on segment-level net sales.	Business diversification

Notes: All variables are constructed using fiscal year 1990 Compustat data. Firm-level financial ratios and diversification measures are aggregated to the CBSA level using employment weights. The fourth column outlines the specific economic dimension of firm quality and structural resilience captured by each indicator.

D EVENT STUDY

D.1 Alternative Specification: Detrended Wages

D.1.1 City Size and Shock Absorption

This section complements Section 4 of the main paper with an alternative specification that removes city-specific pre-trends in log wages prior to PNTR and run the regression using de-trended log wages as the dependent variable. Specifically, for each city c, we first estimate a linear city-specific pre-trend using pre-shock observations,

$$\ln w_{ct} = \alpha_c + \rho_c t + u_{ct}, \qquad t \leq 1999, \tag{D.18}$$

and construct the detrended outcome for all years as

$$\widetilde{\ln w}_{ct} \equiv \ln w_{ct} - \widehat{\rho}_c t. \tag{D.19}$$

We then replace $\ln w_{ct}$ with $\widetilde{\ln w}_{ct}$ in Eq. 8 and re-estimate the specification with the same set of controls and fixed effects. Figure **??** plots the resulting estimates, indicating that the baseline findings are robust to accounting for city-specific pre-trends in wages.

REFERENCES

Acemoglu, D., Carvalho, V. M., Ozdaglar, A., and Tahbaz-Salehi, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016.

Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The china syndrome: Local labor market effects of import competition in the united states. *American economic review*, 103(6):2121–2168.

Carvalho, V. M. (2014). From micro to macro via production networks. *Journal of Economic Perspectives*, 28(4):23–48.

Conley, T. G. and Dupor, B. (2003). A spatial analysis of sectoral complementarity. *Journal of political Economy*, 111(2):311–352.

Davis, J. C. and Henderson, J. V. (2008). The agglomeration of headquarters. *Regional science and urban economics*, 38(5):445–460.

Dupor, B. (1999). Aggregation and irrelevance in multi-sector models. *Journal of Monetary Economics*, 43(2):391–409.

Kopytov, A., Taschereau-Dumouchel, M., and Xu, Z. (2025). The origin of risk. Working paper.

Shea, J. (2002). Complementarities and comovements. *Journal of Money, Credit & Banking*, 34(2):412–434.