INDUSTRIAL CONCENTRATION, PROPERTY VALUES, AND MUNICIPAL BOND SPREADS*

KENNETH R. AHERN

Abstract

This paper proposes that the composition of a city's industrial base helps determine the pricing of its bonds through the interaction of two channels. First, cities that are home to a diversified set of industries can better absorb shocks to any single sector. Second, cities dominated by high-wage industries are more affluent and have greater fiscal capacity to absorb shocks. Using a panel of 1,253 cities from 2005 to 2022, I show that employment concentration raises municipal bond spreads—but only in cities with low property values. A two-channel instrumental variables design using trade-exposure shocks and housing supply elasticities provides evidence that concentration and property values each causally affect spreads. A calibrated model of stochastic industrial dynamics shows that the two channels are inherently linked: under heterogeneous wages across sectors, the same growth process that concentrates employment also determines property values. The model implies that fiscally vulnerable cities can mobilize roughly one-quarter the emergency fiscal capacity of wealthy cities, amplifying the borrowing cost of concentration in less affluent cities. These results help explain the wide dispersion in borrowing costs across the $4 trillion U.S. municipal bond market.

This Version: March 15, 2026

JEL Classification: G12, H74, R51

Keywords: Municipal bonds, credit spreads, industrial diversification, fiscal risk, property values

The author did not receive any source of research support external to the University of Southern California and he has no potential conflicts of interest.

* University of Southern California and NBER. E-mail: kenneth.ahern@marshall.usc.edu.

I. Introduction

The U.S. municipal bond market accounts for roughly $4 trillion of securities issued by more than 50,000 state and local government entities, making it one of the most important, but fragmented credit markets in the world.[1] Though municipal bonds issued by city governments are often depicted as homogenous, low-risk assets, protected by state and federal backstops, the reality is that many issuers face severe fiscal distress including the well known bankruptcies of San Bernardino, California, and Detroit, Michigan. Moreover, fiscal distress is increasing. Since 2021, the growth in the average city's expenses has outpaced the growth in its revenues (National League of Cities, 2025). This shift is reflected in the rising number of cities declaring fiscal emergencies, including a wide range of cities, from Los Angeles to Albany, New York.

The borrowing costs of city governments have important real-world consequences. Higher spreads translate directly into higher infrastructure costs, reduced public investment, and ultimately lower-quality public services for residents. For the average person, the most tangible services provided by any government entity are provided by city governments, including schools, roadways, sidewalks, police, and fire departments. Yet, the existing literature on the determinants of municipal bond spreads is small and there is much that is unknown.

In this paper, I propose that a first-order determinant of a city bond's yield spread is the composition of industries in its local economy. First, cities that are supporting by a diverse set of industries can more easily absorb negative shocks from any single industry without catastrophic revenue loss. This follows from the fundamental tenet of finance that diversification reduces total volatility and hence expected distress costs (Merton, 1974; Leland, 1994). Second, a city's composition of industries is a primary determinant of its affluence and affluent cities have latent fiscal buffers unavailable to poorer city that they can draw on when shocks occur. This simple framework is supported by survey responses from city finance officers who identify the value of the local tax base (68%) and the health of the local economy (50%) as the two most important determinants of their fiscal outlook—well ahead of state support (25%) (National League of Cities, 2025).

To illustrate my hypothesis, consider Gary, Indiana (population 86,000) and Hammond, Indiana (population 78,000), two neighboring cities in the Calumet industrial corridor along the southern shore of Lake Michigan. Both cities share the same state-level tax and institutional environment, similar demographic profiles, and a common history of steel and heavy

[1]Securities Industry and Financial Markets Association (SIFMA), U.S. Municipal Bonds Statistics; U.S. Securities and Exchange Commission, Office of Municipal Securities.

manufacturing. In recent years, Gary's economy remains heavily concentrated in manufacturing, with nearly 25% of employment in a single sector. Hammond, by contrast, has diversified into healthcare, retail, education, and professional services, with manufacturing accounting for only 12% of employment. When the steel industry contracted—driven by import competition, automation, and the 2008–2009 recession—Gary's narrow economic base amplified the fiscal shock, with property values falling to a median of $68,000, and its municipal bond spreads averaging 230 basis points over the 2005–2022 sample period, peaking at 420 basis points in 2011. Hammond, exposed to the same regional steel decline, absorbed the shock through its diversified employment base. Its property values remained higher ($96,000 median), and its bond spreads averaged only 106 basis points—less than half of Gary's. The contrast between these two neighboring cities, sharing a common state and regional economy but differing in industrial concentration and fiscal capacity, reflects the central hypothesis of this paper.

To test my hypothesis, I construct a city-year panel of 1,253 cities from 2005 to 2022 by matching secondary-market municipal bond trading data from the MSRB to employment data from the U.S. Census Bureau's LODES program. The employment data provide sector-level job counts at the city level, from which I compute the Herfindahl–Hirschman Index (HHI) of employment concentration across 20 NAICS supersectors. The bond data provide par-value-weighted average yields, from which I compute maturity-matched credit spreads over the Treasury curve.

In the first set of tests, I isolate the effect of HHI on yield spreads by controlling for city and year fixed effects plus a host of time-varying explanatory variables, including liquidity measures and city population. The results show that employment concentration predicts higher bond spreads, but the relationship is concentrated among cities with limited fiscal capacity. Among cities in the lowest tercile of median property values, a one-standard-deviation increase in HHI is associated with a 5.6 basis point increase in maturity-matched spreads, relative to a mean of 29 bps. This is a sizable economic magnitude, equivalent to the effect of a one-standard-deviation increase in the unemployment rate. Among cities in the highest tercile of property values, the effect is economically and statistically indistinguishable from zero.

Though the fixed effects and time-varying controls address the most likely confounding variables, unobserved omitted variables could explain the result. Because HHI and property values are jointly determined by local economic conditions, neither can be treated as exogenous. To address this concern, I develop a two-channel instrumental variables design that jointly instruments employment concentration and property values. HHI is instrumented by

a Bartik shift-share predictor that interacts predetermined city-industry employment shares with exogenous national sector trends derived from TiVA trade data, following Goldsmith-Pinkham, Sorkin, and Swift (2020) and Borusyak, Hull, and Jaravel (2022). Log property value is instrumented by a cumulative housing shock that interacts city-specific housing supply elasticities with regional house-price movements (Guren, McKay, Nakamura, and Steinsson, 2021). Flow validation tests confirm that each instrument predicts its target endogenous variable but not the other, supporting the channel separation required for identification. The two-stage least squares results confirm the OLS findings: increases in HHI cause higher yield spreads, and higher property values lower spreads. The magnitude of the IV estimates are larger than the OLS estimates, consistent with attenuation bias from measurement error in the 20-sector HHI. Though the two-channel design identifies the separate causal effects of concentration and property values, I do not estimate a causal interaction because the same economic process that drives industrial concentration also determines property values through sector-specific wages. In other words, there is no third source of exogenous variation that moves the interaction independently of its components. Instead, the structural model developed below resolves this limitation.

I also consider three alternative explanations. First, municipal bonds are disproportionately held by local retail investors (Coval and Moskowitz, 1999; Schultz, 2012), and concentrated cities could exhibit stronger local bias. If local investors are overoptimistic, this would push yields *down* and make my estimates a lower bound; if local ownership reduces liquidity, it could inflate spreads. I show that the HHI coefficient is stable—and if anything increases—when augmented with four bond-level liquidity measures, ruling out the liquidity channel. Second, concentrated cities might simply have weaker governance; city fixed effects and the Bartik instrument, which isolates variation driven by foreign demand shocks, address this concern. Third, state-level institutional constraints such as balanced-budget requirements and bankruptcy authorization are time-invariant and absorbed by city fixed effects; the results are further robust to state $\times$ year fixed effects.

To provide a structural interpretation of the paper's main results, I develop a calibrated model of stochastic industrial dynamics and fiscal default. The model rests on two building blocks. First, a city's industrial composition evolves according to a Gibrat process with mean reversion and entry. This single stochastic process jointly determines both a city's employment concentration and, through heterogeneous sector wages, its property values—revealing that the two channels identified separately in the IV design are structurally entangled. Second, property values are endogenous to sector composition (Rosen, 1979; Roback, 1982): cities dominated by high-wage industries have high property values because those sectors

generate the wages and amenities that are capitalized into housing prices. The combination of heterogenous wages across sectors and a Gibrat growth process allows the model to closely match both the HHI and the variation in property values observed in the data.

In the model, cities raise revenue from property taxes and non-property tax sources, pay for government services, save a fiscal surplus, and service debt. Bondholders price municipal debts based on the risk neutral probability of default. Default occurs when a cumulative employment-driven revenue shortfall exceeds the city's fiscal buffer, which scales with property wealth. The model produces a key structural parameter, ψ, that quantifies a city's latent fiscal capacity—the fraction of its property wealth that can be mobilized as an emergency fiscal buffer. The calibrated value $\psi^* = 0.380$ implies that cities can tap approximately 4% of their normalized property wealth per year through assessment increases, millage adjustments, and rainy-day fund drawdowns. Because high-property-value cities have property wealth that is nearly four times as large as low property-value cities, the calibrated model matches the flat HHI–spread gradient for wealthy cities and the steep gradient for fiscally vulnerable cities. Thus, the model provides a structural explanation for the interaction between concentration and property values documented in the reduced-form evidence: property wealth acts as a fiscal buffer that insulates wealthy cities from the revenue volatility generated by concentrated employment.

The central contribution of this paper is to identify the industrial composition of a city's employment base as a first-order determinant of its municipal borrowing costs. While the existing literature has examined other determinants of municipal bond pricing—including credit ratings (Cornaggia, Cornaggia, and Israelsen, 2018; Adelino, Cunha, and Ferreira, 2017), tax policy (Longstaff, 2011; Poterba, 1989; Babina, Jotikasthira, Lundblad, and Ramadorai, 2021), pension liabilities (Novy-Marx and Rauh, 2012), political institutions (Butler, Fauver, and Mortal, 2009), state distress policies (Gao, Lee, and Murphy, 2019), local information environments (Gao, Lee, and Murphy, 2020), climate risk (Goldsmith-Pinkham, Gustafson, Lewis, and Schwert, 2022), demographic change (Butler, Fauver, and Mortal, 2022), and place-based subsidies (Chava, Malakar, and Singh, 2024)—no prior work has studied how a city's industrial diversification affects its cost of debt. This gap is surprising given that a city's industrial mix is arguably the most fundamental structural characteristic of its local economy, shaping both the level and volatility of the tax base.

Unlike many previously studied determinants (e.g., ratings recalibrations, policy changes, or information shocks), industrial composition is a slow-moving, deeply structural feature that is difficult for cities to change and difficult for researchers to identify in reduced-form regressions. The existing municipal bond literature—including the foundational work of

Schwert (2017), Green, Hollifield, and Schürhoff (2007), and Novy-Marx and Rauh (2012)—has largely documented cross-sectional correlations without formal causal identification. The combination of a Bartik IV strategy, a two-channel design that separates concentration from fiscal capacity, and a calibrated structural model that reveals their structural entanglement goes substantially beyond this standard—not because OLS evidence is uninformative, but because the slow-moving nature of industrial composition makes a stronger design desirable.

II. Data and Sample Construction

I construct a city-year panel of municipal bond spreads and employment concentration by linking secondary-market bond trading records to Census employment data. This section describes the data sources, variable construction, and sample composition.

II.A. Municipal Bond Trading Data

Bond pricing data come from the Municipal Securities Rulemaking Board (MSRB) via its Electronic Municipal Market Access (EMMA) system. The MSRB records every municipal bond trade executed in the United States, including trade date, yield-to-maturity, par value, and counterparty type. I restrict to customer trades, excluding inter-dealer transactions that represent inventory transfers between dealers rather than arm's-length market prices (Harris and Piwowar, 2006; Schwert, 2017). I compute the maturity-matched spread for each individual trade (equation 1) and then aggregate to the city-year level using par-value weights, producing a value-weighted average spread for each city in each year from 2005 to 2022. I require at least three trades per city-year to form an observation, reducing noise from thinly traded issues. The final sample includes cities from 43 states plus Washington D.C.[2]

I compute the maturity-matched bond spread by interpolating each trade's yield to the corresponding constant-maturity Treasury point on the yield curve:

$$\text{Spread}_{ct} \;=\; \frac{1}{\sum_j \text{par}_j} \sum_j \text{par}_j \left[y_j - r^{\text{CMT}}(\text{mat}_j) \right] \times 100, \tag{1}$$

where y_j is the yield-to-maturity of trade j, $r^{\text{CMT}}(\text{mat}_j)$ is the constant-maturity Treasury yield interpolated to the remaining maturity of trade j, and the spread is expressed in basis

[2]States with small populations tend to have thinly traded bonds. Hence, the sample does not include any cities in Alaska, Delaware, Hawaii, Idaho, North Dakota, and South Dakota. In addition, the state-level Local Government Commission of North Carolina sells all municipal bonds on behalf of local governments which prevents a match between the MSRB data and the Census data.

points. The Treasury benchmark is constructed from the Federal Reserve's H.15 constant-maturity Treasury (CMT) series at seven tenors (1, 2, 5, 7, 10, 20, and 30 years), downloaded from FRED. Bonds with remaining maturity below six months are excluded to avoid near-maturity noise; bonds with remaining maturity above 30 years are benchmarked against the 30-year CMT rate. The maturity-matched spread accounts for the term structure of Treasury yields rather than benchmarking all trades against a single maturity point, producing a more precise measure of municipal credit risk. As a robustness check, I also compute a tax-adjusted spread following Schwert (2017) and find similar qualitative results. All spread and yield variables are winsorized at the within-year 1st and 99th percentiles.

II.B. Entity Matching

Linking MSRB bond records to Census economic data requires matching MSRB issuer identifiers to Census place or county subdivision codes. I construct this crosswalk in two steps. First, I parse MSRB security descriptions to extract the city name and state associated with each six-digit CUSIP issuer code. Second, I fuzzy-match extracted city names to a master entity concordance of Census places and county subdivisions using token-set-ratio string similarity, requiring a minimum match score of 85 on a 0–100 scale. I manually verify ambiguous matches and correct errors in the automated crosswalk.

For cities in Minor Civil Division (MCD) states—including Connecticut, Massachusetts, New York, New Jersey, Pennsylvania, Ohio, Michigan, and Illinois—local governments are organized as county subdivisions (towns and townships) rather than incorporated places. I use the ACS county subdivision geographic identifier for these entities, ensuring consistent matching across state governance structures.

II.C. Employment Concentration

An important explanatory variable of this paper is the industrial concentration of cities. Because my unit of analysis is a city, it is important that my measure of industrial concentration is at the city level. Using county-level sector distributions would generate imprecise measures of a city's economic condition, especially for smaller cities in large counties. The most precise data available is employment data from the Workplace Area Characteristics (WAC) files of the U.S. Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) Origin-Destination Employment Statistics (LODES).

LODES provides a near-census of jobs covered by state unemployment insurance programs, geocoded to Census blocks and classified into 20 NAICS supersectors (CNS01 through CNS20 in the LEHD classification). The 20-sector aggregation level is coarser than the four-digit

NAICS classifications available from some establishment-level datasets but offers two advantages: LODES provides comprehensive coverage of virtually all employers at the city level, and the supersector classification is consistent across all years and geographies in the sample.[3] I aggregate employment from block level to Census places using the Census Bureau's geographic crosswalk, producing sector-level employment counts for each city and year from 2002 to 2022.

I measure industrial concentration using the Herfindahl–Hirschman Index:

$$\text{HHI}_{ct} = \sum_{k=1}^{K} p_{kct}^2, \tag{2}$$

where $p_{kct} = E_{kct}/E_{ct}$ is sector k's share of total employment in city c in year t, and $K = 20$ is the number of NAICS supersectors. HHI ranges from $1/K = 0.05$ for a perfectly diversified city to 1 for a single-sector economy.

The distribution of sector employment shares within cities is right-skewed: most sectors account for small fractions of local employment, while a few dominant sectors hold disproportionately large shares. This pattern mirrors the well-documented skewness of firm size distributions (Gibrat, 1931; Gabaix, 1999) and establishment counts (Luttmer, 2007), extending those regularities to the within-city sector level. The right skew implies that HHI is driven primarily by the one or two largest sectors in each city, making the concentration measure sensitive to shocks in dominant industries.

The 20-sector NAICS supersector classification is coarser than finer industry classifications available from establishment-level datasets. I use this aggregation for three reasons. First, LODES is the only public dataset that reports employment at the city (place) level rather than at the county or metropolitan-area level. Because municipal debt is issued and repaid by individual cities, city-level employment is essential for measuring the fiscal exposure of each issuer. No establishment-level dataset (such as the County Business Patterns or QCEW) provides place-level geography, and aggregating county data to cities introduces severe measurement error for the many cities that do not coincide with county boundaries. Second, with $K = 20$ sectors and nearly universal coverage across cities, virtually all variation in HHI reflects the *unevenness* of employment shares rather than the number of active sectors. Ahern, Kong, and Yan (2026) show formally that concentration indexes decompose into a component driven by the count of categories and a component driven by the inequality of shares; when K is large relative to dominance, the count margin contributes little. In my sample, most cities have positive employment in all 20 sectors, so cross-city differences in

[3]LODES covers approximately 95 percent of private employment through state unemployment insurance records, plus federal employment. Self-employment and informal employment are excluded.

HHI are driven almost entirely by how unevenly employment is distributed—precisely the margin relevant for revenue-volatility risk. Third, finer classifications suffer from disclosure suppression: at four-digit or six-digit granularity, many city-industry cells are suppressed in public data, artificially inflating measured HHI. I use workplace-based employment (WAC) throughout the analysis because municipal debt is serviced by the taxing jurisdiction's own revenue base. WAC captures economic activity within city boundaries—the employers whose presence generates commercial property value and tax revenue. When a plant closes, the host city's tax base contracts, making workplace-based concentration the relevant measure of fiscal exposure.

II.D. Latent Fiscal Capacity

A key proposal of this paper is that property values capture residents' implicit fiscal collateral: cities with higher property values have greater latent capacity to sustain property tax increases before triggering resident exit (Tiebout, 1956; Oates, 1969). I proxy for fiscal capacity with the American Community Survey (ACS) estimate of median property value at the city level.[4] For the 2009–2022 period, I use ACS 5-year estimates for places and county subdivisions. For the 2005–2008 period, I extend coverage using ACS 3-year estimates (available for places with population above 20,000), ACS 1-year estimates (places above 65,000), and, for remaining gaps, a county-level proxy based on the 2009 ACS 5-year county estimate. I assign each city-year observation to a property value tercile formed within each year's cross-section of ACS median property values, allowing cities to move between terciles as their property values rise or fall.

I interpret median home value as a proxy for the community's aggregate wealth tied to the jurisdiction, rather than exclusively as a measure of the property tax base. This distinction is important because several states impose binding constraints on property tax assessments or rates (e.g., California's Proposition 13). I argue that market property values remain informative as a fiscal capacity proxy for three reasons. First, even in tax-limited states, high property values correlate with higher incomes, sales tax bases, and fee revenues, providing alternative fiscal instruments that cities can mobilize during downturns. Second, property tax limitations typically constrain the *rate of growth* rather than permanently fixing revenues; reassessment upon sale and new construction gradually align assessments with market values. Third, city fixed effects absorb the time-invariant component of state tax regimes, and the state $\times$ year fixed effects robustness check in Appendix C absorbs any time-varying state-level policy changes, including modifications to assessment rules.

[4]ACS table B25077, variable B25077_001E.

II.E. Additional Controls

In all of the regressions, I control for city fixed effects to absorb city-level time-invariant characteristics that could affect municipal bond pricing, such as persistent government structures, legal statutes, and geographic-based traits. I also control for year fixed effects to absorb all macroeconomic conditions common to all cities that could affect the time-series of bond yields.

To control for time-varying factors, I control for log median property value, the county-level unemployment rate (Bureau of Labor Statistics, Local Area Unemployment Statistics), the state top marginal income tax rate (Tax Foundation), and log population (Census Bureau's Population Estimates Program (PEP)).[5] I also control for security-level characteristics scraped from EMMA's security detail pages, including par-weighted average years to maturity to capture the term premium (Schwert, 2017); the city's log number of trades and log par value traded to control for secondary-market liquidity (Harris and Piwowar, 2006);[6] the city's par-weighted callable share of bonds to control for the embedded option premium in callable bonds; and the city's par-weighted general obligation (GO) share to account for the lower spreads earned by bonds that are backed by the full faith and credit of the issuer vis-à-vis revenue bonds which are backed by specific project cash flows (Capeci, 1994).

II.F. Sample Characteristics

The final analysis panel contains 22,762 city-year observations spanning 1,253 cities from 2005 to 2022. The panel is unbalanced: not all cities have sufficient secondary-market trading activity in every year. Of these, 19,402 observations (85 percent) have non-missing employment concentration data from LODES; the remainder are primarily county or special-district issuers for which Census place-level employment data are unavailable. Of the 19,402 observations with HHI data, 18,868 have non-missing values for all regression controls. Of these, 18,858 remain after dropping singleton city-year groups in the fixed effects estimation and constitute the estimation sample reported in the OLS tables.[7]

---

[5]PEP estimates are based on the decennial census and updated annually using administrative records on births, deaths, and migration. For city-years where PEP data are missing from the Census of Governments extract, I supplement with direct PEP vintage files.

[6]Appendix D reports robustness tests using four additional trade-level liquidity measures—bid-ask proxy, yield dispersion, zero-trading-day fraction, and turnover—constructed from the underlying MSRB transactions. The HHI coefficient is robust to all specifications.

[7]The state $\times$ year fixed effects specification (Appendix C) drops an additional 35 singletons, yielding $N =$ 18,823. The IV specifications use a smaller sample ($N =$ 17,557) because the Bartik instrument requires non-missing 2002 base-period employment shares, which excludes cities incorporated or annexed after 2002.

Figure I maps the sample cities, with color indicating average HHI (darker shading corresponds to higher employment concentration). Many cities tend to cluster around large metropolitan areas, though there are many cities that are in more rural areas as well. The level of industrial concentration varies considerably, even in relatively small geographic areas.

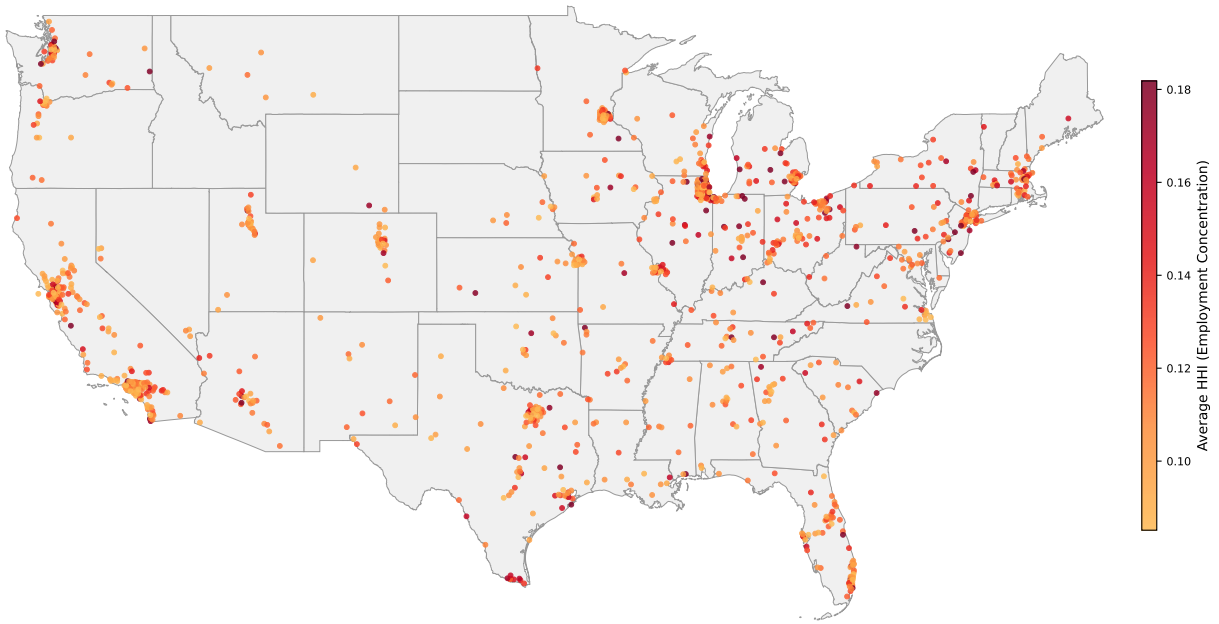


FIGURE I
Geographic Distribution of Sample Cities

Notes: Each dot represents a city in the analysis sample. Darker shading indicates higher employment concentration (average HHI over the 2005–2022 sample period). Larger cities are plotted on top of smaller cities so that metropolitan centers remain visible. The sample spans 1,253 cities across 43 states and Washington D.C.

Table I reports summary statistics for the 18,868 city-year observations with non-missing maturity-matched spread, HHI, and all regression controls. The mean maturity-matched spread is approximately 29 basis points with a standard deviation of 86 basis points. The mean HHI of 0.121 corresponds to an effective number of equal-sized sectors of approximately 8.3, indicating moderate concentration relative to the 20-sector maximum. The median city has a population of approximately 55,000.

These summary statistics are broadly consistent with prior municipal bond studies. Schwert (2017) reports a mean tax-adjusted spread of approximately 106 basis points and a mean remaining maturity of 13 years in a sample of primary-market offerings. The lower mean spread (29 basis points) reflects the maturity-matched benchmark, which interpolates the

Treasury curve at the exact bond maturity rather than benchmarking against a single 10-year Treasury rate with a 40% tax adjustment. Computing the same tax-adjusted spread used by Schwert (2017) in this sample yields a mean of 102 basis points, comparable to his 106 basis points. Gao et al. (2020) report a mean GO share of approximately 50 percent and a median population of roughly 50,000 in their city-level sample, comparable to the GO share and city size distributions in this sample. These statistics indicate the sample is a good representation of the municipal bond market.

TABLE I
Summary Statistics

	Mean	Std. Dev.	P25	Median	P75
Maturity-matched spread (bps)	29.0	85.6	-23.3	19.4	64.8
HHI (employment concentration)	0.121	0.039	0.098	0.111	0.131
Median property value ($000s)	261.8	205.8	132.1	199.7	316.8
GO bond share	0.474	0.397	0.023	0.444	0.915
Average years to maturity	9.8	5.0	6.3	9.1	12.5
Population (000s)	113.7	333.3	36.3	55.2	95.3
Unemployment rate (%)	6.0	2.6	4.1	5.3	7.5
Trades per city-year	786	5,761	63	182	505
Par value traded ($M)	170.2	1787.9	5.4	17.8	53.9

Notes: The sample contains $N = 18{,}868$ city-year observations with non-missing maturity-matched spread, HHI, and all regression controls. The maturity-matched spread is computed by interpolating each trade's yield to the corresponding constant-maturity Treasury point. HHI is the Herfindahl–Hirschman Index of employment concentration across 20 NAICS supersectors from LODES. Median property value is from the American Community Survey. All spread and yield variables are winsorized at the within-year 1st and 99th percentiles.

In unreported tests, I find that larger cities are more diversified, as expected. In particular, industrial concentration is negatively correlated with log population ($\rho \approx -0.30$). I also find that concentration is negatively correlated with log median property value ($\rho \approx -0.15$): less affluent cities tend to be more concentrated.

To illustrate, high-concentration, low-property-value cities include manufacturing-dependent communities such as Decatur, IL and Midland, TX, while low-concentration, high-property-value cities include diversified centers such as Santa Clara, CA and Arlington, VA. The off-diagonal cases—Stamford, CT (concentrated in financial services but high-value) and Dayton, OH (diversified but economically distressed)—confirm that HHI and property values capture distinct dimensions of a city's economic structure, not simply city size.

III. Reduced-Form Empirical Evidence

This section documents that employment concentration predicts higher municipal bond spreads, but that this relationship is concentrated among cities with limited fiscal capacity. I first show univariate evidence that yield spreads are higher for concentrated cities, particularly among cities with low property values, and confirm in Appendix A that the cross-sectional relationship holds in a pure between-city regression with state fixed effects. I then estimate panel OLS regressions with city and year fixed effects and a comprehensive set of controls that isolate within-city variation, progressively tightening the identification to move toward a causal interpretation.

III.A. Univariate Evidence

Table II provides univariate evidence on yield spreads by industrial concentration and property values. Among cities in the lowest tercile of property values, those in the highest tercile of industrial concentration have maturity-matched yield spreads of 34.49 bps compared to 24.04 bps for cities in the lowest tercile, a statistically significant difference. In contrast, among cities with the highest property values, the yield spreads are statistically equal between cities with high concentration (38.08 bps) and low concentration (34.55 bps). These results provide the first evidence that industrial concentration is correlated with higher costs of debt, but only when the latent fiscal capacity of the city is low. Notably, mean spreads are higher for high-property-value cities across all concentration categories, likely reflecting composition effects: these cities tend to issue bonds with longer maturities and more complex structures. The multivariate regressions in Table III, which control for maturity, callable share, and other bond characteristics, reverse this pattern.

Table IA.1 in Appendix A confirms that the univariate patterns survive multivariate controls in a pure cross-sectional setting. Collapsing the panel to city-level means and regressing average spreads on average HHI with state fixed effects, the baseline HHI coefficient is 59.8 basis points but imprecisely estimated (SE 45.8). However, interacting HHI with property value terciles reveals a clear gradient: among low-home-value cities, the HHI coefficient is 96.0 basis points ($p < 0.10$), declining to 77.9 for middle-tercile cities and an insignificant 10.4 for high-value cities. This cross-sectional evidence establishes that the concentration–spread relationship is not an artifact of panel structure or within-city time-series variation.

TABLE II
Mean Bond Spreads by Industrial Concentration and Property Value

	Low Property Value	High Property Value	Difference (High − Low)
High Concentration	34.49	38.08	3.59
			(0.65)
Low Concentration	24.04	34.55	10.51**
			(2.31)
Difference (High − Low)	10.45**	3.53	-6.92
	(2.33)	(0.63)	(-0.96)

Notes: Entries are mean maturity-matched bond spreads in basis points, averaged at the city level over the sample period. Cities are split into terciles (High/Low Concentration) and into property-value terciles (High/Low Property Value). Cell counts: High Concentration/Low Property Value = 228, High Concentration/High Property Value = 168, Low Concentration/Low Property Value = 161, Low Concentration/High Property Value = 220. Middle terciles are omitted. Robust t-statistics in parentheses. Difference-in-differences estimated by OLS with robust standard errors. ***, **, * indicate significance at the 1%, 5%, and 10% levels, respectively.

III.B. Multivariate Fixed Effects Regressions

To isolate the effect of employment concentration from potential confounders, Table III reports panel OLS estimates with city and year fixed effects that absorb all time-invariant city characteristics and common time trends, and a comprehensive set of controls. Standard errors are two-way clustered by city and year. The city fixed effects hold constant any time-invariant trait such as state-level laws, geographic assets, and infrastructure, which helps to ensure that identification comes from within-city changes in HHI over time, rather than from cross-sectional differences across cities. For completeness, I also present purely cross-sectional tests in Appendix A.

Column (1) shows that after including controls, the HHI coefficient is 54.0 basis points (SE 39.8; $t = 1.36$, $p = 0.19$), positive but not statistically significant at conventional levels. However, as hypothesized, Column (2) presents results consistent with the interaction of concentration and property values. Among low-fiscal-capacity cities, the HHI coefficient is 143.6 basis points (SE 52.8; $t = 2.72$, $p = 0.015$), statistically significant at the 5 percent level: a one-standard-deviation increase in HHI is associated with a 5.6 basis point increase in maturity-matched spreads for these cities. The interaction terms show that the gradient diminishes significantly for wealthier cities. The HHI×medium property value interaction is −133.3 basis points (SE 49.1; $p = 0.015$), implying a net gradient of approximately 10 basis points for middle-tercile cities. The HHI×high property value interaction is −148.6 basis

points (SE 62.4; $p = 0.029$), implying a net gradient of approximately -5 basis points for high-property-value cities—effectively zero. The difference between the highest and lowest property value terciles of 149 basis points is statistically significant at the 5 percent level, confirming that the concentration premium is concentrated among cities with limited fiscal capacity.

To contextualize the magnitudes, consider a low-property-value city at the 75th percentile of HHI (0.131) relative to the 25th percentile (0.098). The coefficient of 143.6 basis points on the low tercile of property value implies a spread difference of $143.6 \times (0.131 - 0.098) \approx 4.7$ basis points for this interquartile move, roughly 16% of the mean spread. For the median low-property-value city in the sample with approximately $50 million of outstanding debt, a 5.6 basis point spread increase (one standard deviation of HHI) implies approximately $28,000 per year in additional interest payments, or roughly $210,000 over a 10-year bond life (undiscounted)—a meaningful fraction of the city's annual fiscal surplus. This magnitude is comparable to estimates from prior research. Schwert (2017) estimates that default risk accounts for 74 to 84 percent of the average tax-adjusted municipal bond spread. In this sample, the mean tax-adjusted spread is 102 basis points—comparable to the 106 basis points reported by Schwert (2017)—implying a default-risk component of approximately 76–86 basis points. Gao et al. (2020) find that newspaper closures—which reduce external monitoring of local government finances—raise municipal bond offering yields by 5–11 basis points. The one-standard-deviation effect of 5.6 basis points for low-property-value cities is of comparable magnitude to the newspaper-closure effect, indicating that employment concentration is an economically meaningful determinant of borrowing costs for fiscally vulnerable municipalities. For high-property-value cities, the concentration premium is approximately zero in the OLS: diversified fiscal capacity substantially moderates the employment-concentration channel.

These results are robust to—and in fact strengthen under—the inclusion of state $\times$ year fixed effects, which absorb all time-varying state-level policies including changes to fiscal oversight, tax regimes, and macroeconomic conditions. Appendix Table IA.3 reports that the low-property-value HHI gradient increases from 143.6 to 170.7 basis points ($p < 0.01$) under this more demanding specification, while the interaction terms remain highly significant. The fact that the results strengthen when state-level confounders are absorbed provides additional confidence that the HHI–spread relationship reflects city-level industrial structure rather than state-level factors correlated with concentration.

TABLE III
Industrial Concentration and Municipal Bond Spreads: OLS Evidence

	(1)	(2)
Industrial Concentration	54.0	143.6**
	(39.8)	(52.8)
Ind. Concentration × Medium Property Value		-133.3**
		(49.1)
Ind. Concentration × High Property Value		-148.6**
		(62.4)
Medium Property Value (tercile)		25.8***
		(6.4)
High Property Value (tercile)		32.2***
		(8.4)
Log Median Property Value	-25.8***	-32.2***
	(8.0)	(8.8)
GO Bond Share	-31.2***	-30.9***
	(5.8)	(5.8)
Callable Share	5.0	5.2
	(12.3)	(12.3)
Average Years to Maturity	2.5***	2.5***
	(0.4)	(0.4)
State Income Tax Rate	-4.3**	-4.1**
	(1.9)	(1.9)
Unemployment Rate	6.5***	6.5***
	(1.4)	(1.4)
Log Number of Trades	7.7**	7.6**
	(2.6)	(2.7)
Log Population	-16.2	-16.1
	(12.7)	(12.4)
Log Par Value Traded	-9.6***	-9.6***
	(2.2)	(2.2)
Observations	18858	18858
Fixed effects	City, Year	City, Year
Within R^2	0.070	0.072

Notes: The dependent variable is the maturity-matched bond yield spread in basis points, computed by interpolating each trade's yield to the corresponding constant-maturity Treasury point. All specifications include city and year fixed effects. Standard errors (in parentheses) are two-way clustered by city and year. Column (1) reports the full-sample coefficient. Column (2) interacts Industrial Concentration with annual property value tercile (base category = lowest third; terciles formed within each year's cross-section). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

IV. Causal Evidence

The OLS results establish that employment concentration predicts higher municipal bond spreads, particularly among cities with limited fiscal capacity. However, the OLS gradient may not identify a causal effect if time-varying omitted variables drive both HHI changes and spread movements. Two endogeneity concerns are especially pressing. First, a city-specific economic shock—such as a plant closure or industry decline—could simultaneously concentrate employment (raising HHI) and deteriorate the fiscal outlook (raising spreads), so that HHI is a symptom of the shock rather than a causal driver of credit risk. Second, property values are not merely predetermined city characteristics. They are endogenous responses to local economic conditions and may themselves respond to the same concentration shocks or to the broader local economic environment that concentration affects. Because HHI and property values are jointly determined by local conditions, neither can be treated as exogenous in a regression of spreads on the other, and both must be instrumented simultaneously.

Furthermore, testing the interaction between concentration and property values in a causal framework is both economically unrealistic and practically challenging. The fundamental problem is that both factors are economically entangled. As the structural model in Section V demonstrates, a single stochastic process—Gibrat-law industrial dynamics—simultaneously generates cross-sectional variation in HHI and in property values. When a dominant employer contracts, both the city's employment concentration and its housing demand move together; the interaction of HHI and property values is not a separable third channel but an inherent feature of how local economic shocks propagate through the industrial base and the housing market. An instrument that shifts the product HHI $\times$ ln(PV) independently of its two components would require a source of variation that affects the joint distribution of concentration and property values without operating through either margin separately—a condition that is difficult to satisfy even in principle.

I therefore estimate a two-channel instrumental variables design that jointly instruments employment concentration and property values using two excluded instruments, each targeting one endogenous variable. The identification strategy requires instruments that generate separate, independent sources of variation in HHI and property values—a Bartik shift-share instrument for concentration and an instrument for property values based on Guren et al. (2021) (GMNS). If these results support the causal role of each factor, then the interaction follows from the economic structure of the problem: if concentration increases fiscal risk by amplifying revenue volatility, then the fiscal buffer available to absorb that volatility must moderate the effect. The alternative hypothesis of no interaction effect would require that the mechanism through which concentration affects credit risk is entirely independent of a

city's ability to respond to revenue shocks, which contradicts both the fiscal structure of municipal default and the patterns in the data.

IV.A. Identification Strategy

The empirical model of the relationship between yield spreads and local economic conditions is

$$\text{spread}_{ct} \;=\; \beta_1 \,\text{HHI}_{ct} \;+\; \beta_2 \,\ln(\text{PV}_{ct}) \;+\; \mathbf{X}'_{ct}\boldsymbol{\gamma} \;+\; \alpha_c \;+\; \delta_t \;+\; \varepsilon_{ct}, \tag{3}$$

where HHI_{ct} is city c's employment Herfindahl–Hirschman index in year t, $\ln(\text{PV}_{ct})$ is the log of median property value, $\mathbf{X}_{ct}$ is a vector of controls, and α_c and δ_t are city and year fixed effects. Both HHI and ln(PV) are treated as endogenous. The two excluded instruments are a Bartik predicted HHI measure for the concentration channel and a cumulative housing shock for the property-value channel. The system is exactly identified: two endogenous variables, two excluded instruments.

The coefficient β_1 measures the effect of a unit increase in HHI on the maturity-matched yield spread, holding property values constant. The coefficient β_2 measures the effect of a one-unit increase in log property values—approximately a 100 percent increase—on spreads, holding concentration constant. Jointly estimating both channels is important: omitting either variable from the instrumented system would force the excluded channel's effect into the error term and bias the remaining coefficient.

IV.B. Bartik-Style Instrument for HHI

I construct a shift-share instrument for employment concentration following the Bartik (1991) framework, building on the analysis of regional labor demand shocks developed by Blanchard and Katz (1992). The instrument combines predetermined city-level employment shares with exogenous national-level industry shocks to generate predicted variation in each city's industrial concentration.

Two recent frameworks clarify the identifying assumptions in shift-share designs. Goldsmith-Pinkham et al. (2020) show that shift-share instruments can be interpreted as exposure-weighted averages of sector-specific instruments, with identification requiring exogeneity of the base-period exposure shares: each city's initial industry mix must be uncorrelated with unobserved determinants of spreads, conditional on controls. In my setting, this condition is unlikely to be met. Local industrial structure reflects prior development patterns, geographic advantages, amenities, and institutional choices, all of which may also affect borrowing costs through channels other than concentration. While I can use shares that

are predetermined—fixed before the bond-spread sample begins—they are not randomly assigned, and the Goldsmith-Pinkham et al. (2020) identifying assumption may not be fully credible.

Borusyak et al. (2022) provide an alternative framework that shifts the identification burden from the shares to the shocks. Under their approach, validity requires that the industry-level shocks are as-good-as-randomly assigned conditional on controls, while the exposure shares serve only as weights that transmit the shocks to local outcomes. This framework is attractive because it does not require that predetermined city-industry shares are exogenous—only that the national shocks are orthogonal to city-level unobservables.

Following the logic of Borusyak et al. (2022), I use industry shocks derived from foreign demand rather than realized domestic employment growth. Specifically, I use the OECD Trade in Value Added (TiVA) database, which reports foreign final demand for US domestic value added (FFD_DVA) by industry and year. These shocks measure shifts in foreign countries' demand for US industry output—variation driven by trading partners' economic conditions, trade policies, and global supply chain realignments, rather than by conditions in any US city. For each industry k and year t, the shock is

$$\tilde{g}_{kt} \;=\; \big[\ln(\mathrm{FFD}_{kt}) - \ln(\mathrm{FFD}_{k,0})\big] \;-\; \big[\ln(\bar{\mathrm{FFD}}_t) - \ln(\bar{\mathrm{FFD}}_0)\big], \tag{4}$$

where FFD_{kt} is foreign final demand for US value added in industry k, the subscript 0 denotes the base year, and $\bar{\mathrm{FFD}}_t = \sum_k \omega_k \mathrm{FFD}_{kt}$ is the employment-weighted aggregate. The subtraction of aggregate foreign demand growth removes the common component absorbed by year fixed effects, leaving only relative cross-industry demand shifts.

Because HHI is a nonlinear function of employment shares—specifically, the sum of squared shares—the instrument cannot be constructed as a standard linear Bartik. Instead, I propagate the exogenous foreign demand shocks through predetermined city employment to construct predicted sectoral employment:

$$\hat{E}_{kct} \;=\; E_{kc,0} \times \exp(\tilde{g}_{kt}), \tag{5}$$

where $E_{kc,0}$ is city c's employment in sector k in the base year. The predicted concentration measure is then

$$\widehat{\mathrm{HHI}}_{ct} \;=\; \sum_{k=1}^{20} \hat{s}_{kct}^2, \qquad \text{where} \quad \hat{s}_{kct} = \frac{\hat{E}_{kct}}{\sum_\ell \hat{E}_{\ell ct}}. \tag{6}$$

This is analogous in spirit to the shift-share designs used by Autor, Dorn, and Hanson (2013) and others, but the aggregation is nonlinear because the target variable is concentration rather than employment growth. The instrument remains a valid shift-share construction: it

combines predetermined local exposure (base-year employment $E_{kc,0}$) with exogenous shocks ($\tilde{g}_{kt}$), mapped through the same functional form as the endogenous variable.

The base-year employment data come from the Longitudinal Employer-Household Dynamics (LODES) Workplace Area Characteristics in 2002, three years before the start of the bond-spread sample in 2005. The predetermined shares are fixed before the sample begins and therefore cannot be driven by contemporaneous credit deterioration or fiscal distress.

IV.C. Housing Cycle Instrument for Property Value

Current property values are endogenous to local economic conditions: a negative employment shock that concentrates the industrial base also depresses local housing demand and prices. Simply controlling for property values in the spread regression would not resolve this simultaneity, and instrumenting HHI alone while treating property values as exogenous would bias the concentration estimate if unobserved local shocks affect both HHI and property values.

I construct a housing instrument following the approach developed by Guren et al. (2021), who estimate city-specific sensitivities to regional house-price cycles. The instrument exploits the fact that cities within the same Census region are exposed to common regional housing market forces, but respond with different elasticities due to differences in housing supply constraints, land-use regulation, and geographic features. For city c in Census region r, the annual Guren et al. shock is

$$z_{ct}^{\text{GMNS}} \;=\; \hat{\gamma}_c \times \Delta \ln(\text{HPI}_{rt}), \tag{7}$$

where $\hat{\gamma}_c$ is city c's estimated sensitivity to regional house-price movements and $\Delta \ln(\text{HPI}_{rt})$ is the log change in the regional (Census division) house-price index. The sensitivity $\hat{\gamma}_c$ is estimated from a first-stage regression of city-level house-price growth on regional house-price growth, using data from the pre-sample period; cities with inelastic housing supply (e.g., geographically constrained coastal cities) have larger $\hat{\gamma}_c$, meaning that a given regional house-price shock translates into a larger local price change (Guren et al., 2021).

The annual Guren shock is naturally a flow instrument: it predicts house-price *growth* in a given year. However, the second-stage specification uses log property values in *levels*, not changes. Using a flow instrument for a stock variable creates a mismatch: a shock to house-price growth in year t affects the level of property values in all subsequent years, and a single year's flow carries limited information about the cumulative stock position. I therefore

construct a cumulative instrument that aggregates the annual shocks over time:

$$Z_{ct}^{\text{GMNS}} \;=\; \hat{\gamma}_c \times \sum_{\tau=1}^{t} \Delta \ln(\text{HPI}_{r\tau}), \tag{8}$$

where $\sum_{\tau=1}^{t} \Delta \ln(\text{HPI}_{r\tau})$ is the cumulative log change in the regional house-price index from the base year through year t. Note the distinction from equation (7): the annual instrument z_{ct}^{GMNS} uses the single-period change $\Delta \ln(\text{HPI}_{rt})$, while the cumulative instrument Z_{ct}^{GMNS} uses the accumulated level $\ln(\text{HPI}_{rt})$, capturing the total stock of house-price movements experienced by the city up to year t. The cumulative instrument is the appropriate mapping from a stock-form excluded variable to a stock-form endogenous variable: it captures the accumulated effect of regional housing market conditions on the level of city property values, weighted by the city's structural sensitivity.

IV.D. Main IV Results

Table IV reports the two-channel two-stage least squares (2SLS) estimates. The four columns present a nested sequence of controls: column (1) includes only city and year fixed effects; column (2) adds the unemployment rate and log population; column (3) adds the state top income tax rate; and column (4) further adds bond characteristics and liquidity controls (GO share, callable share, average maturity, log trades, and log par traded). I present this nested sequence for transparency, because some controls may themselves respond to the instruments and thus qualify as "bad controls" (Angrist and Pischke, 2009). The unemployment rate and log population in column (2) are plausibly predetermined within a year, and the state income tax rate in column (3) is set at the state level and unlikely to respond to city-level employment shocks. Column (3) is therefore my preferred specification. The bond characteristics in column (4)—GO share, callable share, average maturity—may partially reflect issuer choices that respond to local economic conditions, creating a potential post-treatment bias. Reassuringly, the HHI and property value coefficients are stable across columns (3) and (4), suggesting that any such bias is small. All specifications use two-way clustering by city and year.

Across all four columns, instrumented HHI enters positively and instrumented log property values enter negatively, consistent with the hypothesis that employment concentration raises borrowing costs while fiscal capacity lowers them. The HHI coefficient declines monotonically as controls are added—from 1,836 basis points with fixed effects only to 832 basis points with the full control set—consistent with the controls absorbing omitted variation that is

TABLE IV
Two-Channel IV: Industrial Concentration and Housing Wealth

	(1)	(2)	(3)	(4)
Industrial Concentration (HHI)	1836.4**	1075.8*	930.5**	832.1**
	(827.6)	(548.6)	(387.9)	(385.2)
	[18.4]	[10.8]	[9.3]	[8.3]
log(Median Home Value)	-335.5**	-285.6**	-157.0**	-159.3**
	(131.1)	(105.9)	(57.1)	(57.4)
	[-3.4]	[-2.9]	[-1.6]	[-1.6]
Unemployment Rate		7.3**	7.0***	6.5***
		(2.5)	(1.3)	(1.2)
log(Population)		147.4**	70.8*	71.3*
		(67.7)	(36.0)	(35.6)
GO Bond Share				-30.9***
				(6.6)
Callable Share				11.6
				(14.9)
Avg. Maturity (years)				2.2***
				(0.4)
log(N Trades)				7.3**
				(2.8)
log(Par Traded)				-8.7***
				(2.4)
State Income Tax Rate (%)			-7.7***	-6.4***
			(1.9)	(1.9)
Observations	16679	16367	16367	16273
Fixed effects	City, Year	City, Year	City, Year	City, Year
Clustering	City, Year	City, Year	City, Year	City, Year
Kleibergen-Paap F	6.6	6.8	13.7	13.9

Dependent variable: maturity-matched yield spread (bps).
HHI instrumented by TiVA Bartik predicted HHI; log(home value) instrumented by cumulative Guren HP shock.
Standard errors two-way clustered by city and year in parentheses.
Bracketed values report implied spread changes (bps) from a 1 percentage point increase in HHI and a 1 percent increase in median home value.
Column (1) includes only city and year fixed effects.
Column (2) adds unemployment rate and log(population).
Column (3) adds state income tax rate.
Column (4) adds bond characteristics and liquidity controls.

positively correlated with both HHI and spreads. The sign and significance of the HHI coefficient remain stable across specifications.

The preferred specification is column (3), which includes unemployment, population, and the state income tax rate. The HHI coefficient is 931 basis points (SE 388; $p < 0.05$), implying that a one percentage point increase in HHI raises the maturity-matched yield spread by approximately 9.3 basis points. The property-value coefficient is −157 basis points (SE 57; $p < 0.02$), implying that a one percent increase in median property value lowers spreads by approximately 1.6 basis points. The Kleibergen–Paap Wald F-statistic is 13.7, comfortably above the Stock and Yogo (2005) 10% maximal IV size critical value of 7.03.

Column (4) adds bond-level and liquidity controls. The HHI coefficient declines modestly to 832 basis points (SE 385), and the property-value coefficient is essentially unchanged at −159 basis points (SE 57). The stability of the estimates when adding bond and liquidity characteristics indicates that secondary-market illiquidity is not driving the concentration premium.

The state income tax rate enters significantly and absorbs part of the property-value effect between columns (2) and (3): the property-value coefficient attenuates from −286 to −157 basis points. This is economically sensible. The tax-exempt status of municipal bond interest means that the state marginal tax rate directly affects the pricing of municipal yields relative to taxable benchmarks. Including the state tax rate sharpens the identification of both channels and substantially strengthens the first stage (the Kleibergen–Paap F-statistic increases from 6.8 to 13.7), motivating its inclusion in the preferred specification.

The IV point estimates are substantially larger than OLS. This is not uncommon in Bartik-style designs (Autor et al., 2013). The IV identifies a local average treatment effect (LATE) for complier cities whose concentration responds most to national sector trends—plausibly smaller cities with tradable-sector exposure, where concentration changes have outsized fiscal consequences. The OLS interaction specification partitions the gradient across property-value terciles, while the IV estimates a pooled linear effect across the full distribution; the unconditional IV coefficient captures a weighted average that gives more weight to the steeper portion of the gradient where the LATE concentrates.

IV.E. First-Stage Evidence and Channel Separation

In a system with multiple endogenous variables, one cannot evaluate instrument strength by examining a single first-stage F-statistic in isolation. The relevant question is whether the instruments jointly provide sufficient and distinct sources of variation for each endogenous regressor—that is, whether each instrument primarily moves its intended endogenous

variable without substantially contaminating the other channel. The first-stage regressions, underidentification tests, and weak-IV-robust diagnostics are all informative about this question.

Table V reports the first-stage regressions. In the preferred specification with controls, the Bartik-predicted HHI instrument enters strongly in the HHI first stage (coefficient 0.184, SE 0.054, $t = 3.41$) and the cumulative GMNS shock enters strongly in the log property-value first stage (coefficient 0.153, SE 0.053, $t = 2.90$). Critically, the cross-predictions are small: the Guren shock has no predictive power for HHI (coefficient -0.003, SE 0.002, $t = -1.39$), confirming that housing market shocks do not affect employment concentration. The Bartik instrument has some correlation with log property values in levels (coefficient 0.753, SE 0.399, $t = 1.89$), which is not surprising: national industry shocks that shift a city's employment base can also affect local property values through labor demand and income channels. However, this correlation is modest and attenuated relative to the specification without controls, suggesting that the unemployment rate and population controls absorb the direct employment-to-housing channel.

This pattern of channel separation—each instrument primarily moving its own endogenous variable—is exactly what is required for the two-channel design to be informative. If both instruments moved both endogenous variables proportionally, the system would be effectively identified by a single common shock, and the separate β_1 and β_2 coefficients would not be distinguishable. The low cross-prediction of GMNS on HHI, combined with the strong own-channel predictions, provides confidence that the two instruments generate distinct sources of variation.

IV.F. Weak-IV Diagnostics

Appendix Table IA.6 reports identification and weak-instrument diagnostics for the preferred specification. The Kleibergen–Paap rank LM statistic of 7.96 ($p = 0.005$) rejects the null of underidentification, confirming that the instruments provide sufficient rank condition for the two-endogenous-variable system. The Kleibergen–Paap Wald F-statistic of 13.7 exceeds the Stock and Yogo (2005) 10% maximal IV size critical value of 7.03, substantially mitigating concerns about weak-instrument bias.

As a further check, I compute a weak-IV-robust joint test of the null hypothesis that both structural coefficients are zero ($\beta_1 = \beta_2 = 0$). The Anderson–Rubin F-statistic of 3.72 ($p = 0.046$) rejects at the 5% level (Staiger and Stock, 1997). Unlike standard Wald tests, the Anderson–Rubin test has correct size regardless of instrument strength, providing an

TABLE V
First-Stage Regressions

	HHI		ln(PV)	
	(1)	(2)	(3)	(4)
Bartik Predicted HHI	0.175***	0.184***	0.888**	0.753*
	(0.054)	(0.054)	(0.394)	(0.399)
Cumulative Guren HP Shock	−0.000	−0.003	0.118**	0.152***
	(0.002)	(0.002)	(0.049)	(0.053)
Unemployment Rate		−0.000		−0.000
		(0.000)		(0.008)
log(Population)		−0.021***		0.532***
		(0.005)		(0.080)
State Income Tax Rate (%)		0.001**		−0.016*
		(0.000)		(0.008)
Within R^2	0.013	0.025	0.014	0.082
Partial F	10.7	11.6	5.9	8.4
Fixed effects	City, Year	City, Year	City, Year	City, Year
Observations	16,679	16,367	16,679	16,367

Notes: The dependent variable in columns (1)–(2) is the employment HHI and in columns (3)–(4) is ln(PV), the log of median property value. Columns (1) and (3) include only city and year fixed effects. Columns (2) and (4) add controls. Bartik Predicted HHI is the shift-share instrument using national industry trends (TiVA trade shocks) applied to local employment shares. Cumulative Guren HP Shock is the city-specific housing supply elasticity times cumulative regional house price changes. Standard errors two-way clustered by city and year in parentheses. Partial F reports the Kleibergen-Paap first-stage F-statistic for each endogenous variable. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

inference procedure that does not rely on the instruments being strong. These diagnostics collectively indicate that the instruments provide meaningful identifying variation.

IV.G. Flow Validation

A natural concern with the cumulative GMNS instrument is whether the stock-form aggregation introduces spurious correlation. The annual GMNS shock $z_{ct}^{\text{GMNS}} = \hat{\gamma}_c \times \Delta \ln(\text{HPI}_{rt})$ is a flow instrument that directly predicts house-price growth. If the instrument is valid, it should predict annual changes in property values but not annual changes in HHI; conversely, the change in Bartik-predicted HHI should predict annual HHI changes but not property-value changes.

Appendix Table IA.5 reports first-difference regressions that confirm this pattern. The annual GMNS shock strongly predicts changes in log property values (coefficient 0.135,

SE 0.015, $t = 8.91$) and has zero predictive power for HHI changes (coefficient 0.000, SE 0.003). The change in Bartik predicted HHI predicts HHI changes (coefficient 0.054, SE 0.023, $t = 2.39$) and has no correlation with property-value changes (coefficient -0.000, SE 0.064). The channel separation is clean in first differences, validating the instrument construction and confirming that the modest Bartik–property-value correlation in the levels first stage reflects accumulated general-equilibrium effects rather than a violation of the exclusion restriction.

The concentration channel is weaker in annual first differences ($t = 2.39$) than in levels ($t = 3.41$) because HHI is a slow-moving variable: year-to-year changes in predicted concentration are small relative to the cumulative cross-sectional variation that the levels specification exploits. The levels panel is therefore the preferred specification for the joint two-channel design, with the first-difference results serving as a validation exercise rather than the primary identification.

IV.H. Alternative Explanations

Before turning to the structural model, I consider several alternative explanations for the observed HHI–spread relationship.

H.1. Local bias and liquidity

Municipal bonds are disproportionately held by local retail investors, a well-documented pattern of home bias in fixed income markets (Coval and Moskowitz, 1999; Schultz, 2012). For example, Babcock, Cummins, and Parsons (2019) show that bonds issued by historically Black colleges and universities (HBCUs) trade at lower yields than comparable non-HBCU bonds, consistent with local investor clienteles accepting below-market returns out of affinity or familiarity. If investors in concentrated cities exhibit stronger local bias—perhaps because a dominant employer creates a more cohesive local identity—two opposing forces could be at work. On one hand, overoptimistic local investors would bid up prices and push yields *down*, which would attenuate the HHI–spread relationship and make my estimates a lower bound on the true credit-risk effect. On the other hand, concentrated local ownership could reduce secondary-market liquidity, inflating spreads through a liquidity premium rather than a default premium.

To distinguish these channels, Appendix D augments the baseline specification with four bond-level liquidity measures constructed from the trade data: a bid-ask proxy (Harris and Piwowar, 2006), yield dispersion (Green et al., 2007), the zero-trading-day fraction (Lesmond, Ogden, and Trzcinka, 1999), and turnover. The HHI coefficient is stable—and if anything

increases—when these controls are added, ruling out liquidity as the primary channel and suggesting that any local-bias effect works in the opposite direction, making the reported estimates conservative.

H.2. Management quality.

Concentrated cities might have weaker governance or less sophisticated fiscal management, and it is this mismanagement—rather than industrial structure per se—that raises borrowing costs. While management quality is difficult to observe directly (Bloom and Van Reenen, 2007), several features of the research design address this concern. City fixed effects absorb persistent differences in governance quality across cities. Time-varying controls for population and unemployment capture changes in the scale of city operations. Most importantly, the Bartik instrument isolates variation in HHI driven by foreign demand shocks to national industries, which are plausibly orthogonal to changes in local management quality.

H.3. Institutional constraints.

Most states require municipalities to maintain balanced budgets, and only about half of states authorize municipal Chapter 9 bankruptcy (Gao et al., 2019). These institutional features could affect the relationship between fiscal stress and borrowing costs. However, state bankruptcy authorization is time-invariant and is absorbed by city fixed effects. Appendix C further shows that the results are robust to the inclusion of state $\times$ year fixed effects, which absorb all time-varying state-level policies including changes to fiscal oversight regimes.

V. Quantitative Model

The empirical regressions in Sections III and IV establish that industrial concentration raises bond spreads, and that this effect is concentrated among cities with low property values. This section develops a quantitative model to rationalize these estimates and to reveal a structural insight that the reduced-form analysis cannot: employment concentration and property values are not independent city characteristics but joint outcomes of the same underlying stochastic process—Gibrat-law industrial dynamics. A single set of sector shocks simultaneously determines a city's HHI and its property values, which is precisely why the interaction of these two variables cannot be cleanly identified with separate instruments (Section IV). In particular, the model matches the key asymmetry in the reduced form regressions—a steep HHI–spread gradient for low-property-value cities and a near-zero gradient for high-property-value cities—with only two free parameters.

The model builds on prior work on municipal credit risk—including Epple and Spatt (1986) on debt limits as commitment devices, Capeci (1994) on borrowing-cost elasticities, and Schwert (2017) on default-liquidity decomposition—but differs in two key respects: default risk arises endogenously from the interaction of stochastic industrial dynamics with the city's fiscal structure, and property values are endogenous to sector composition, creating a structural link between industrial mix, fiscal capacity, and credit spreads.

V.A. Model Overview

The economy consists of a large number of cities, each populated by a set of industrial sectors that evolve stochastically. Sectors grow and decline according to a Gibrat process—proportional random growth—which generates the skewed, approximately lognormal size distributions observed in firm sizes (Gibrat, 1931), city populations (Gabaix, 1999), and establishment counts (Luttmer, 2007). In a standard Gibrat process, the variance of the firm size distribution grows without bound. To generate a stationary distribution, I follow the existing literature and allow for mean reversion (Gabaix, 1999) and sector entry (Luttmer, 2007), both standard features of applied Gibrat models.

Cities collect taxes and service debt according to their observed fiscal structure, and bondholders price municipal debt at risk-neutral spreads implied by default probabilities. Though the model abstracts from agents that maximize a city's objective function, the fiscal ratios observed in the data—expenditure shares, leverage ratios, property tax reliance—are treated as the outcome of cities' optimal fiscal policy and taken as inputs rather than derived endogenously.[8] The model targets the isolated marginal effects of HHI and property value estimated from regressions that include city and year fixed effects and a full set of controls. By matching to these marginal effects, the model does not need to explain all of the confounding variables, only the marginal effect of HHI and property value.

The estimation proceeds in two stages for computational tractability. Stage 1 calibrates the stochastic process governing industrial dynamics to match the cross-sectional distribution of employment concentration (HHI) observed in the data, simultaneously generating an endogenous distribution of property values from sector composition. This joint determination is the structural foundation for the entanglement of HHI and property values discussed in Section IV. Stage 2 takes the simulated industrial structure as given and calibrates the fiscal default mechanism to match three bond-spread moments from the regression estimates. Two

[8]In principle, these fiscal ratios could be microfounded through a Buchanan club model in which residents choose tax and expenditure levels. Such an extension would not change the model's predictions about the HHI–spread gradient, since the key mechanism operates through the interaction of employment shocks with the existing fiscal structure, not through the optimality of that structure.

parameters are calibrated: the baseline default probability attributable to all non-HHI factors and the latent fiscal capacity that generates the gradient of spreads with respect to HHI, conditional on property value.

V.B. Stage 1: Industrial Dynamics and Endogenous Property Values

The first stage of the model generates the distribution of sector sizes within each city and the cross-sectional heterogeneity in property values across cities. Each city begins with K_0 active sectors drawn randomly from J total sectors in the economy. Because sector identities determine property values, sector identities are not interchangeable, as they are in generic Gibrat models.

In each of $T_{\text{burn}} = 200$ periods, active sectors experience log-normal employment shocks:

$$\ln E_{k,t+1} = \ln E_{kt} + \sigma_{\text{eff}} \cdot \varepsilon_{kt}, \qquad \varepsilon_{kt} \sim \mathcal{N}(0,1), \tag{9}$$

where $\sigma_{\text{eff}} = \sigma(1 - \rho \cdot p_{kt})$ is a size-dependent effective volatility.[9] In generic Gibrat models, the variance of the distribution of sector sizes grows each period and the distribution of sector sizes is non-stationary. To overcome this limitation, I modify the Gibrat process following the prior literature on stochastic growth models (de Wit, 2005). First, the mean-reversion parameter ρ dampens volatility for large sectors, preventing any single sector from dominating the city. Second, sectors with employment shares below a minimum threshold exit, while new sectors enter with probability λ per period at a small initial share p_0. The city's industrial concentration is $\text{HHI}_c = \sum_k p_{ck}^2$ as in equation (2).

The same process that determines industrial concentration also determines a city's property values. Following a standard urban economics model (Rosen, 1979; Roback, 1982), a city's property values are determined by its industrial composition, where sectors with higher wages increase property values. Cities dominated by high-premium sectors (e.g., technology, finance) have higher property values; cities dominated by low-premium sectors (e.g., manufacturing, extraction) have lower property values. Specifically, the log of median property value is a share-weighted sum of sector-specific property value premia v_k, reflecting the capitalized contribution of each industry to local economic conditions:

$$\ln(\text{PV}_c) = \alpha + \sum_{k=1}^{J} p_{ck} \cdot v_k. \tag{10}$$

[9]The use of independent Normal shocks across sectors is conservative: if sectors were positively correlated (e.g., during recessions), diversification would provide less protection, and the HHI gradient would be steeper than what the model generates.

Thus, in the steady state, each city's employment shares $\mathbf{p}_c = (p_{c1}, \ldots, p_{cJ})$ determine both its HHI and its property value.

V.C. Stage 2: Fiscal Structure and Default

In the steady state of industry composition and property values, cities receive revenues, pay expenses, and face random shocks which cause bond yield spreads to be determined by the endogenous probability of default.

C.1. City Revenues and Expenses

Each city c collects revenue from two sources: property taxes (share ω_c of total revenue) and employment-sensitive revenue (share $1 - \omega_c$), which includes income taxes, sales taxes, service fees, and intergovernmental transfers. Data from the Census of Governments (U.S. Census Bureau, 2017; Ahern, 2024) show that cities with higher property values derive a larger share of revenue from property taxes, consistent with less elastic property bases supporting greater reliance on this revenue source. The property tax share of revenues is allowed to vary across cities to reflect optimal taxation strategies. Cities with relatively less elastic property values optimally rely more heavily on property taxes to raise revenue. Pre-shock total revenue is normalized to unity to control for variation in city size.

Cities have two types of expenses: non-debt expenditure g_c (as a fraction of revenue) and debt service $d_c \cdot (r_f + s_c)$, where d_c is the leverage ratio (long-term debt to revenue), r_f is the risk-free rate, and s_c is the credit spread. Non-debt expenditure and the leverage ratio are exogenous and constant across cities. This restriction of the model is supported by empirical evidence and institutional constraints. In the Census of Governments data, expenditure-to-revenue ratios are remarkably stable across property value terciles (≈ 0.94) (Ahern, 2024), and leverage ratios are similarly constrained by state-imposed debt limits and balanced-budget requirements that restrict municipal fiscal flexibility (Kiewiet and Szakaly, 1996). In contrast, empirical credit spreads vary significantly across cities. Therefore, in the model, the credit spread, s_c, is heterogeneous and endogenously determined.

The difference between revenues and expenses, $1 - g_c - d_c(r_f + s_c)$, reflects a city's annual fiscal surplus. Historically, municipalities maintain fund balances to smooth revenue fluctuations and meet unexpected expenditures. The Government Finance Officers Association recommends minimum unrestricted fund balances of at least two months' expenditure, or approximately 17% of annual revenue (Government Finance Officers Association, 2015). This surplus serves as a buffer against negative shocks.

The assumption of exogenous fiscal ratios warrants discussion. In principle, cities could adjust expenditure and leverage in response to employment shocks, and endogenizing this margin would require a full dynamic public finance model. However, the empirical evidence suggests that this margin is limited. State-imposed balanced-budget requirements and debt limitations (Kiewiet and Szakaly, 1996; Epple and Spatt, 1986) severely constrain municipal fiscal flexibility, and the expenditure-to-revenue ratio of approximately 0.94 is remarkably stable across property value terciles. Moreover, the composition of fixed versus variable expenditures (personnel, infrastructure, debt service) is approximately constant across high- and low-property-value cities, so any symmetric ability to cut expenditures would attenuate the HHI–spread gradient equally for all cities, leaving the *differential* between high and low property value cities unchanged. The approach is analogous to structural corporate credit risk models (Merton, 1974; Leland, 1994) that take leverage as given when pricing credit risk, even though leverage is ultimately a choice variable. The model does not endogenize fiscal policy choices; this is a deliberate simplification that isolates the diversification–property value channel from the broader set of fiscal policy responses.

C.2. Employment and property value shocks

Just as the sector-level dynamics of Stage 1 determine both industrial concentration and property values, sector-level shocks drive variance in concentration and property values. In particular, bond investors assess default risk over the remaining T periods in the life of the bond. Following Gibrat dynamics, each sector's cumulative employment change has standard deviation $\sigma\sqrt{T}$. Thus, cumulative sector shocks are

$$g_k \sim \mathcal{N}(0, \sigma^2 T), \qquad k = 1, \ldots, J. \tag{11}$$

City-level employment change is the share-weighted sum of cumulative sector shocks, $\Delta E_c/E_c = \sum_k p_{ck} \cdot g_k$. For a diversified city, positive and negative shocks partially cancel; for a concentrated city, outcomes are dominated by its largest sector.

The same sector shocks that affect employment also affect property values through (10). A city's post-shock log property value is:

$$\ln(\mathrm{PV}_c^{\mathrm{post}}) = \alpha + \sum_k p_{ck}^{\mathrm{post}} \cdot v_k, \tag{12}$$

where the post-shock shares p_{ck} reflect differential sector growth. The ratio $\mathrm{PV}_c^{\mathrm{post}}/\mathrm{PV}_c^{\mathrm{pre}}$ captures the endogenous property value response. A negative shock to any sector reduces property values, but the magnitude depends on which sector contracts: a shock to a high-value sector (e.g., Finance) causes a larger decline than a shock to a low-value sector (e.g.,

Arts). The interaction between heterogeneous sector values and the skewed distribution of sectors creates an asymmetric response between high- and low-value cities for identical changes in concentration.

C.3. Default probabilities and credit spreads

The terminal revenue shortfall at horizon T, as a fraction of initial annual revenue, is:

$$\ell_c^T = -(1-\omega_c)\cdot\frac{\Delta E_c}{E_c} - \omega_c\cdot\left(\frac{\mathrm{PV}_c^{\mathrm{post}}}{\mathrm{PV}_c^{\mathrm{pre}}}-1\right). \tag{13}$$

The first term is the loss of employment-sensitive revenue; the second is the loss of property tax revenue caused by the erosion of the property tax base. Both terms are positive when the city experiences a negative shock.

The structure is analogous to a Merton-style barrier model: sector shocks accumulate as a random walk over T periods, so the terminal draw has variance $\sigma^2 T$ (equation (11)), and default occurs if the realized loss exceeds the city's total absorption capacity:

$$\ell_c^T > \underbrace{T\cdot(1-g_c-d_c(r_f+s_c))}_{\text{cumulative fiscal surplus}} + \underbrace{\psi\cdot\frac{\mathrm{PV}_c}{\overline{\mathrm{PV}}}}_{\text{emergency property wealth}}. \tag{14}$$

The first term on the right is the cumulative fiscal surplus over T years. The second term is the city's latent fiscal capacity: the ability to raise additional revenue by tapping its property wealth through assessment increases, millage adjustments, special assessments, or rainy-day fund drawdowns. The parameter ψ governs the strength of this channel, and $\mathrm{PV}_c/\overline{\mathrm{PV}}$ is the city's property value normalized by the cross-city mean. Thus, the parameter ψ reflects the fraction of a city's excess property wealth that is available to meet excess debt obligations.

The structural default condition given by equation (14) generates a probability of default, q_c^{HHI}, determined by the random sector shocks, which drive both HHI and property-value heterogeneity. In addition, I assume that default can also occur with probability q_{base} determined by non-modeled sources of fiscal distress: pension underfunding, legal liabilities, infrastructure crises, and political dysfunction. Including this residual default probability makes it less likely that heterogeneity in the data caused by factors unrelated to concentration will be artificially attributed to the concentration channel.

Thus, the total default probability is:

$$q_c = q_{\mathrm{base}} + q_c^{\mathrm{HHI}}. \tag{15}$$

Bondholders price municipal debt at the risk-neutral spread:

$$s_c = \frac{(1 + r_f)(1 - \delta)q_c}{1 - (1 - \delta)q_c}, \tag{16}$$

where δ is the recovery rate. Since q_c depends on s_c through (14), the spread is calculated by a fixed-point iteration with 50% damping until convergence.

C.4. The Model's Mechanism

The model generates asymmetry between high- and low-property-value cities through three reinforcing channels, all endogenous to sector composition. No separate insurance or convexity parameter is needed—the protection arises from the economic structure itself.

The first channel is shock diversification. In a diversified city, positive and negative sector shocks partially cancel across the portfolio of sectors. The net employment change has standard deviation proportional to $\sigma\sqrt{T \cdot h_c}$, so a city with HHI = 0.05 has half the employment volatility of a city with HHI = 0.20. This diversification benefit is pure portfolio theory and does not depend on property values.

The second channel is property tax stability. Property values are determined by sector composition through (10). When a low-value sector contracts, high-value cities (which are dominated by high-value sectors) experience a small relative decline in $\text{PV}^{\text{post}}/\text{PV}^{\text{pre}}$. Low-value cities, whose sectors are uniformly low-value, experience proportionally larger declines. Since property tax revenue is proportional to property values, high-value cities maintain a more stable revenue base (Lutz, Molloy, and Shan, 2011).

The third channel is the latent fiscal capacity. The term $\psi \cdot \text{PV}_c/\overline{\text{PV}}$ in (14) means high-value cities have a larger emergency buffer—not because of an arbitrary parameter assignment, but because wealthier cities have greater latent taxing capacity.

V.D. Parameters and Identification

The calibration of the model to the data is completed in two stages. The first stage estimates the parameter values that govern the Gibrat process to match the observed variation in industrial concentration and property values. The second stage estimates the parameter values that govern the relationship between the yield spread and concentration.

D.1. First Stage Estimation

The goal of the first stage is to identify the parameter values that govern sector employment volatility, entry-probability, and mean-reversion. I run 2,000 Monte Carlo simulations of the

Gibrat process across a three-dimensional grid of 945 parameter combinations ($9\times7\times5$ values of σ, λ, and ρ, crossed with three K_0 configurations). Using a weighted loss function, I target five moments of the cross-sectional HHI distribution: mean, standard deviation, skewness, the mean HHI for small cities, and the mean HHI for large cities. This matches the overall shape of the HHI distribution accounting for the fact that small cities are more concentrated than large cities. From these five moments, I recover the following four parameter values: $\sigma = 0.080$ (sector employment volatility), $\lambda = 0.200$ (sector entry probability), $\rho = 0.60$ (mean-reversion strength), and $K_0 \in [6, 12]$ for small cities and $K_0 \in [14, 20]$ for large cities (the initial sector count). Panel A of Table VI reports the full parameter set.

Property values are not targeted in the first-stage calibration. Instead, they are assigned based on the stationary equilibrium and estimated sector loadings. In particular, I estimate equation (10) from a between-city regression of log median property value on sector employment shares using all cities in the calibration panel. This regression estimates sector values $\{v_k\}_{k=1}^{J}$ and intercept α, with an $R^2 = 0.43$, indicating that sector composition alone explains a substantial share of the cross-city variation in property values. The sector values range from $v = 15.4$ (Other Services) and $v = 10.6$ (Real Estate) at the top to $v = -3.0$ (Utilities) and $v = -3.6$ (Mining) at the bottom.

The sector-value coefficients are fixed data inputs, not calibrated parameters, that I then apply to the steady state equilibrium to compare the model distribution of property value to the data. A close match indicates that the calibrated Gibrat process matches both industry concentration and property value heterogeneity, though only industry concentration is used in the calibration.

Table VI compares the model-implied data and the actual data for the first stage calibration. The model matches the HHI distribution well in the body (mean 0.111 vs. 0.120 in data, SD 0.039 vs. 0.039) but understates the extreme right tail (skewness 2.1 vs. 3.6), reflecting a well-known limitation of Gibrat processes relative to the Pareto tails observed in concentration data. To the extent that this understates extreme concentration, the model provides a conservative estimate of the structural spread gradient, since the most concentrated cities drive the largest default probabilities.

Though the model is not explicitly calibrated to property values, the model-generated SD(ln property value) = 0.60 and the mean ln(property value) = 12.22, matching the data exactly, consistent with the economic entanglement of concentration and property values. Figure II plots the joint bivariate distribution of HHI and log property values in the data (left) and model (right). Both distributions share the same support and similar shapes, centered near HHI ≈ 0.10 and ln(PV) ≈ 12.2. The data exhibit a modest negative correlation ($\rho =$

TABLE VI
Stage 1: Gibrat Stochastic Growth Model with Endogenous Home Values

	Model	Data
Panel A: Industry dynamics		
σ (sector volatility)	0.080	—
λ (entry rate)	0.200	—
ρ (mean reversion)	0.60	—
K_0 distribution	K=[6-12]/[14-20]	—
Panel B: Model fit (multi-seed average)		
Mean of HHI	0.1112	0.1200
Standard deviation of HHI	0.0391	0.0393
Skewness of HHI	2.1	3.6
Mean of log(property value)	12.22	12.22
Standard deviation of log(property value)	0.60	0.60
Corr(HHI, ln(property value))	-0.018	-0.094
Panel C: Sector values (from data, $R^2 = 0.43$)		
Other Services	15.4	Fixed
Real Estate	10.6	Fixed
Professional	5.7	Fixed
Agriculture	4.6	Fixed
Information	3.3	Fixed
...		
Arts	-0.3	Fixed
Utilities	-3.0	Fixed
Mining	-3.6	Fixed
Intercept (α)	10.6	Fixed

Notes: Structural fiscal default model with endogenous home values. Standard errors across simulation seeds in parentheses. Each of 20 NAICS sectors has a value v_k estimated from a between-city regression of log(median home value) on sector employment shares ($R^2 = 0.43$). Sector identities are assigned at city initialization and fixed throughout Gibrat dynamics. Each sector receives a Normal$(0, \sigma^2)$ growth shock; positive and negative shocks net out for diversified cities but not for concentrated ones. Panel B averages over 8 MC seeds. Sample: cities with publicly traded GO bonds.

-0.094): more concentrated cities tend to have somewhat lower property values. The model generates a weaker correlation ($\rho = -0.018$) because sector identities are assigned randomly at initialization rather than reflecting geographic or historical sorting. This difference does not affect the spread calibration, which conditions on the property value tercile (and hence on the realized property value) rather than relying on the correlation between HHI and property values.

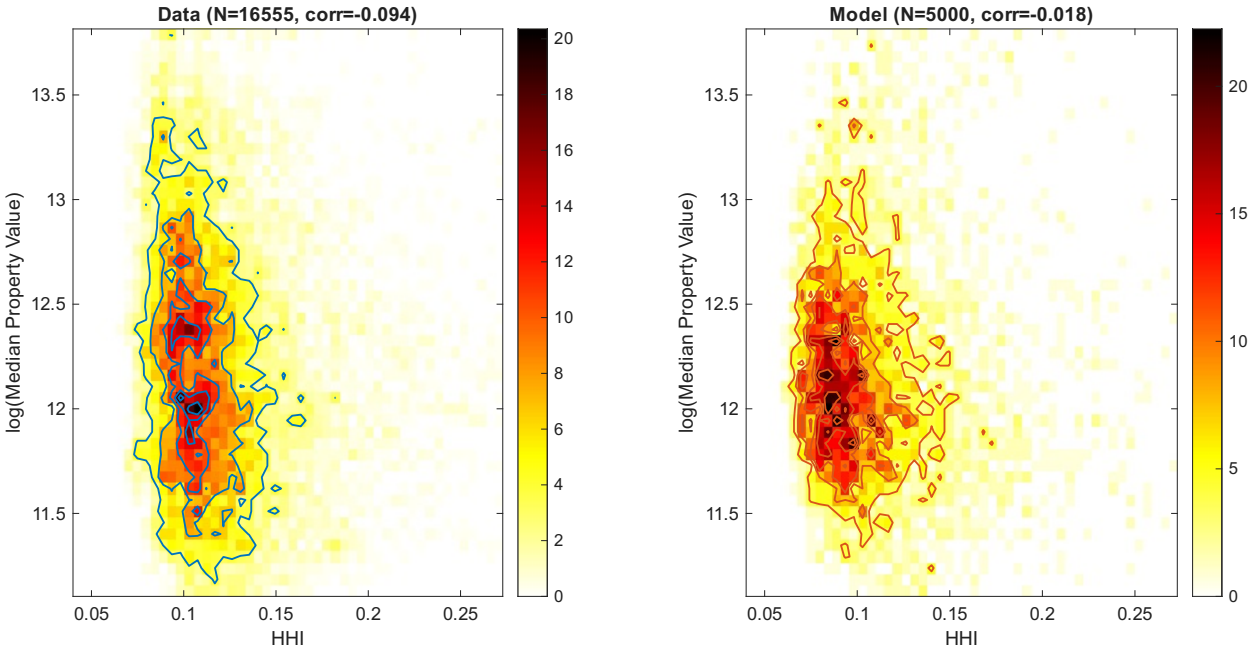


FIGURE II
Joint Distribution of HHI and Property Values: Data vs. Model

Notes: Bivariate heatmap of HHI × log(median property value) in the data (left, $N = 16{,}555$ city-years, $\rho = -0.094$) and model (right, $N = 5{,}000$ cities, $\rho = -0.018$). Color intensity indicates density; contour lines show probability mass. Both distributions share similar support and shape. The data exhibit a modest negative correlation (concentrated cities tend to have lower property values); the model generates a weaker correlation because sector identities are assigned randomly at initialization. Property values are endogenous to sector composition via (10).

D.2. Second Stage Estimation

The second stage calibration uses the first stage sector dynamics results plus a set of data inputs that are fixed externally to identify two free parameters: the latent fiscal capacity ψ and the baseline default probability q_{base}. Panel A of Table VII lists the fixed parameter values taken from the data. Fiscal ratios (ω_c, g_c, d_c) are computed as their average values in the Census of Governments, conditional on the city's property value tercile. Cities in the highest tercile of property values have larger ω_c (0.29 vs. 0.18 for lowest tercile), so more of their revenue comes from the relatively stable property tax base, consistent with less elastic property values. In contrast, high value and low-value cities have similar expenditure ratios (≈ 0.94), and similar leverage ratios (≈ 1.1–1.2). This confirms that cities are likely to have relatively little flexibility to change expenditures and leverage when shocks arise.

The maturity horizon $T = 9.4$ years is the par-value-weighted mean remaining maturity of bonds in the spread regression sample (the equal-weighted mean is 9.8 years; I use the

par-weighted mean because the spread measure is itself par-weighted). Its pricing relevance is confirmed by the regression coefficient of 2.5 bps per year ($t = 6.73$). The risk-free rate $r_f = 0.03$ and recovery rate $\delta = 0.40$ follow standard values in the structural credit risk literature (Leland, 1994; Duffie and Singleton, 1999).

To estimate the structural parameters of the model, I calculate $q_c^{\text{HHI}} = \Pr(\ell_c^T > \text{threshold}_c)$ from 20,000 Monte Carlo draws of sector shocks. Importantly, the four Gibrat parameters $(\sigma, \lambda, \rho, K_0)$ and the 20 sector values $\{v_k\}$ and intercept α from the first stage are used to determine the cumulative shock volatility that drives the default mechanism: $\sigma\sqrt{T} = 0.080 \times \sqrt{9.4} \approx 0.245$. The identical $\sigma = 0.080$ that generates the cross-sectional HHI distribution in Stage 1 also governs the fiscal shocks that drive default in Stage 2. There is no separate "shock process" for pricing—the volatility that produces the observed concentration distribution is the same volatility that determines credit risk. This tight linkage disciplines the model: the data on industrial structure directly constrain the default mechanism.

The system is overidentified: two free parameters target three moments from the empirical results (low-property-value city gradient = 143.6 bps, high-property-value city gradient = 0.0 bps, and mean spread = 29.0 bps). Parameters are estimated by minimizing the sum of squared proportional deviations using a two-dimensional grid search over (ψ, q_{base}). Overidentification is a deliberate feature, not an accident—the model must simultaneously match the three moments, and there is no guarantee that a two-parameter model can do so. If the model can fit the three moments, it indicates that the fiscal-buffer mechanism parsimoniously captures the key economic forces.

Table VII reports the calibrated parameters and model fit. First, the results show that the model provides a close fit to the data. Panel B of Table VII reports the in-sample fit, averaged over 8 Monte Carlo seeds (each with 20,000 shock draws for the same 5,000 cities). The HHI gradient for low-value cities is 135.8 ± 0.6 bps versus the 143.6 bps target (95% of target). The HHI gradient for high-value cities is 0.1 ± 0.0 bps versus 0.0 in the data. The mean spread is 28.8 ± 0.0 bps versus 29.0. All three moments are matched closely, and the simulation variance across seeds is small.

Panel C reports an out-of-sample validation. The in-sample fit (Panel B) holds the 5,000 simulated cities from Stage 1 fixed and draws fresh employment shocks—this tests whether the calibrated parameters produce stable spread moments for the same cities under different shock realizations. The out-of-sample test (Panel C) goes further: each of 8 independent seeds generates 5,000 entirely new cities from scratch (new Gibrat Stage 1 simulations producing new industrial compositions and property values, new fiscal-ratio draws from the Census data) and computes spreads at the calibrated $(\psi^*, q^*_{\text{base}})$. These cities have never been

TABLE VII
Stage 2: Gibrat Model with Endogenous Home Values

	Model	Data
Panel A: Fiscal ratios (from data)		
Low property value cities		
Property tax share (ω_i)	0.173	Fixed
Expenditure ratio (g_i)	0.934	Fixed
Leverage (d_i)	1.185	Fixed
High property value cities		
Property tax share (ω_i)	0.292	Fixed
Expenditure ratio (g_i)	0.942	Fixed
Leverage (d_i)	1.140	Fixed
Maturity horizon, years (T)	9.4	Fixed
Risk-free rate (r_f)	0.03	Fixed
Recovery rate (δ)	0.40	Fixed
Panel B: Model fit (multi-seed average)		
Mean spread (bps)	28.8 (0.0)	29.0
HHI gradient, low property value cities (bps)	135.8 (0.6)	143.6
HHI gradient, high property value cities (bps)	0.1 (0.0)	0.0
Mean default prob (%)	0.46	—
Panel C: Out-of-sample (8 seeds × 5K cities)		
Mean spread (bps)	28.8 (0.1)	29.0
HHI gradient, low property value cities (bps)	148.5 (17.4)	143.6
HHI gradient, high property value cities (bps)	0.1 (0.0)	0.0
Mean default prob (%)	0.46	—
Panel D: Structural default parameters		
Latent fiscal capacity (ψ)	0.380	—
Baseline default (q_{base})	0.0045	—

Notes: Structural fiscal default model with endogenous home values. Standard errors across simulation seeds in parentheses. Each of 20 NAICS sectors has a value v_k estimated from a between-city regression of log(median home value) on sector employment shares ($R^2 = 0.43$). Sector identities are assigned at city initialization and fixed throughout Gibrat dynamics. Each sector receives a Normal$(0, \sigma^2)$ growth shock; positive and negative shocks net out for diversified cities but not for concentrated ones. Cumulative revenue shortfall over T years: $\text{eff_loss}_T = -(1-p)(\Delta E/E) - p(HV_{\text{post}}/HV_{\text{pre}} - 1)$, where shocks scale as $\sigma\sqrt{T}$ over the bond maturity horizon ($T = 9.4$ years from data). Default when $\text{eff_loss}_T > T(1 - g - d(r_f + s)) + \psi \times HV_i/\overline{HV}$, i.e., cumulative loss exceeds T years of fiscal surplus plus emergency property wealth (ψ), or from idiosyncratic fiscal distress (q_{base}). Spreads priced via risk-neutral fixed point: $s_i = (1 + r_f)(1 - \delta)q_i/[1 - (1 - \delta)q_i]$. Only two free parameters in Stage 2: ψ (property wealth fraction) and q_{base}. Fiscal ratios from Census of Governments conditional on property value tercile. Panel C averages over 8 MC seeds; Panel D draws fresh cities each seed. Sample: cities with publicly traded GO bonds.

seen during calibration. The out-of-sample low-property-value gradient is 148.5 ± 17.4 bps, closely centered on the 143.6 bps target. The high-property-value gradient remains flat (0.1 ± 0.0 bps) and the mean spread is 28.8 ± 0.1 bps. The larger out-of-sample standard error for low-property-value cities reflects the sensitivity of the gradient to the realized HHI distribution in each fresh city draw. As an additional validation, the middle-property-value gradient is not targeted in calibration yet provides an untargeted test of the model's predictive power: the empirical HP2 gradient is approximately 10 bps (Table III), and the model implies an intermediate gradient that declines smoothly from the HP1 to HP3 values, consistent with the empirical pattern.

Figure III plots the model-implied spread gradient against the data for each property value tercile. The left panel shows a positive relationship between HHI and spreads in both data and model in low-value cities, with the model line tracking within the 95% confidence intervals of the data. The right panel shows a flat relationship in both data and model for high-value cities.

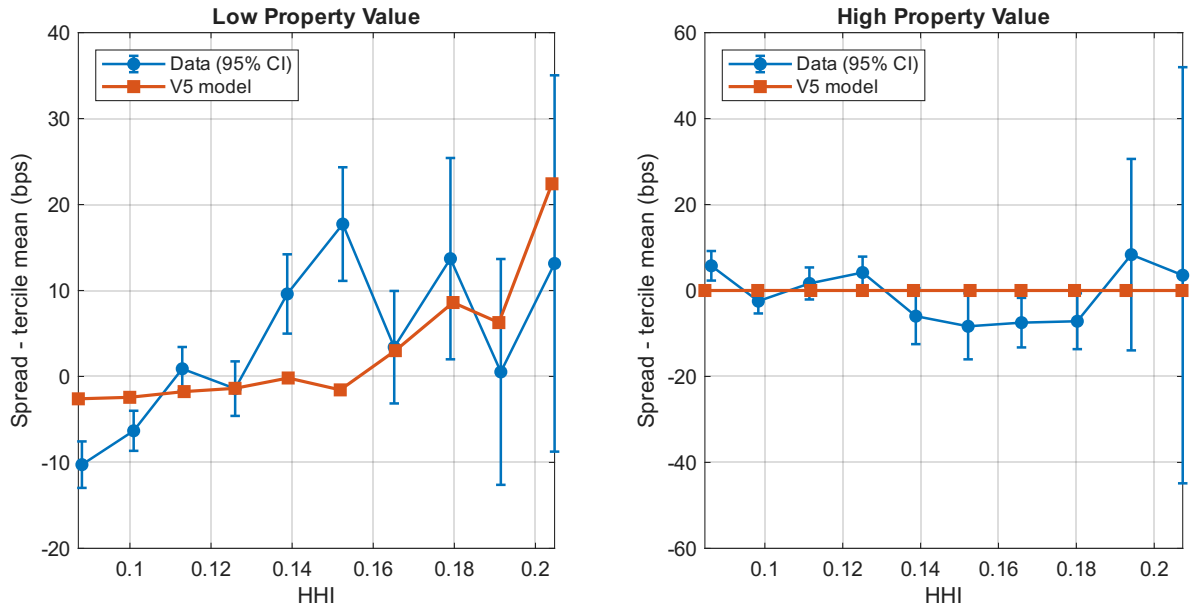


FIGURE III
Model Fit: HHI–Spread Gradient by Property Value Tercile

Notes: Binned scatter of the HHI–spread relationship by property value tercile (low = left, high = right). Data points show the econometric prediction from the city and year fixed effects regression, demeaned within each property value tercile; error bars are 95% confidence intervals. Model values (squares) are from the calibrated simulation at $\psi^* = 0.380$, $q^*_{\text{base}} = 0.0045$, $T = 9.4$ yr. The left panel (low property value) shows a steep positive gradient in both data and model. The right panel (high property value) shows a flat gradient, indicating that high-property-value cities are largely buffered against concentration risk.

Having established that the model fits the data closely, I turn to the calibrated values of the structural parameters: $\psi^* = 0.380$ and $q^*_{\text{base}} = 0.0045$. The calibrated value $\psi^* = 0.380$ implies that cities can mobilize approximately 38% of their normalized property wealth over the $T = 9.4$ year horizon, or approximately 4% per year—consistent with incremental millage increases, special assessments, and rainy-day fund drawdowns. The economic content of ψ is that it translates the industrial composition of a city—through the property value channel—into a fiscal buffer. Cities with high-value sectors (Real Estate, Professional Services) have high $\text{PV}_c/\overline{\text{PV}}$ and therefore large latent capacity; cities with low-value sectors have small capacity. In the data, high-value cities have $\text{PV}_c/\overline{\text{PV}} \approx 1.8$; low-value cities have $\text{PV}_c/\overline{\text{PV}} \approx 0.5$. This nearly 4-to-1 ratio means high-property-value cities can absorb substantially larger cumulative losses before defaulting. This is the structural mechanism by which industry composition determines borrowing costs: rich cities can borrow more cheaply because their industrial base generates property values that provide insurance against fiscal shocks.

The baseline default probability $q^*_{\text{base}} = 0.0045$ plays the role of a level shifter, analogous to a regression intercept. It captures the 0.45% probability of default from non-modeled sources (pension underfunding,[10] legal liabilities, political dysfunction), contributing approximately 28 bps to the mean spread. The structural HHI mechanism adds 0–20 bps on top, concentrated among low-property-value cities with high concentration. This decomposition is consistent with the empirical observation that mean municipal default rates are very low (Ang, Bhansali, and Xing, 2014), while the cross-sectional variation in spreads responds to fundamental fiscal risk, paralleling the income-fluctuation mechanism in sovereign default models (Arellano, 2008).

The model is calibrated to the within-city OLS gradients rather than the IV estimates for two reasons. First, the OLS coefficients conditional on city and year fixed effects represent average partial effects for the full cross-section of cities, which is the natural target for a structural model that simulates a heterogeneous population. The IV estimates identify local average treatment effects for complier cities whose concentration responds to national sector shocks—a different estimand that may not be representative of the broader population. Second, and more fundamentally, the model's key target is the interaction between HHI and property values—the differential gradient across the property-value distribution. As discussed in Section IV, this interaction cannot be causally identified because the same

[10]Pension liability data are not available at the city level for the full sample. City fixed effects absorb persistent differences in pension obligations across cities, and the baseline default probability q_{base} captures the average contribution of pension risk to default. To the extent that pension underfunding is correlated with industrial composition—for example, manufacturing cities with legacy defined-benefit plans—the city fixed effects in the OLS specifications absorb this channel, and the Bartik instrument in the IV specifications isolates foreign-demand-driven variation in HHI that is plausibly orthogonal to pension policy.

stochastic process generates both variables. The model provides the structural framework needed to interpret the interaction: it shows why the gradient must differ across property-value terciles, grounding the OLS pattern in the mechanics of fiscal default rather than relying on a causal estimate that the data cannot deliver.

As discussed above, the model allows for three reinforcing channels: shock diversification, stable property tax base, and latent fiscal capacity. Figure IV decomposes the annual fiscal loss into its components by property value tercile. The employment revenue loss (red bars) is similar across terciles, reflecting that all cities face similar average shocks. The property tax erosion (orange) is negligible because property values are relatively insensitive to single-sector contractions. The annual fiscal surplus (blue) provides the primary buffer ($\sim$9% of revenue for all terciles). The latent fiscal capacity (green) is the key differentiator: cities in the top tercile of property values receive approximately 7.3% annualized latent capacity versus 1.9% for cities in the bottom tercile—a nearly 4-to-1 ratio that mirrors the property value differential. For high-value cities, the additional capacity is nearly as large as the annual fiscal surplus itself, effectively doubling their loss-absorption capacity. For low-value cities, the additional capacity is a small fraction of the surplus, leaving them exposed to tail-risk employment shocks.

VI. Conclusion

This paper identifies the industrial composition of a city's employment base as an important determinant of its municipal borrowing costs, particularly for fiscally vulnerable cities. Cities with concentrated employment pay higher credit spreads, but this premium is concentrated among cities with limited fiscal capacity—those with low property values.

The identification strategy uses a two-channel instrumental variables design that jointly instruments employment concentration and property values. A Bartik shift-share instrument, driven by exogenous national sector trends from TiVA trade data, isolates variation in HHI that is orthogonal to local fiscal conditions. A cumulative housing shock instruments log property values through city-specific housing supply elasticities interacted with regional house-price movements. Flow validation tests confirm that each instrument predicts its target endogenous variable but not the other. The instrumental variables estimates of the HHI effect are larger than OLS, consistent with attenuation bias from measurement error in the 20-sector HHI or a local average treatment effect that upweights cities whose concentration responds most to national sector shocks. The heterogeneity by property values survives—and strengthens—under instrumentation, confirming that the pattern reflects a structural relationship rather than a spurious correlation.

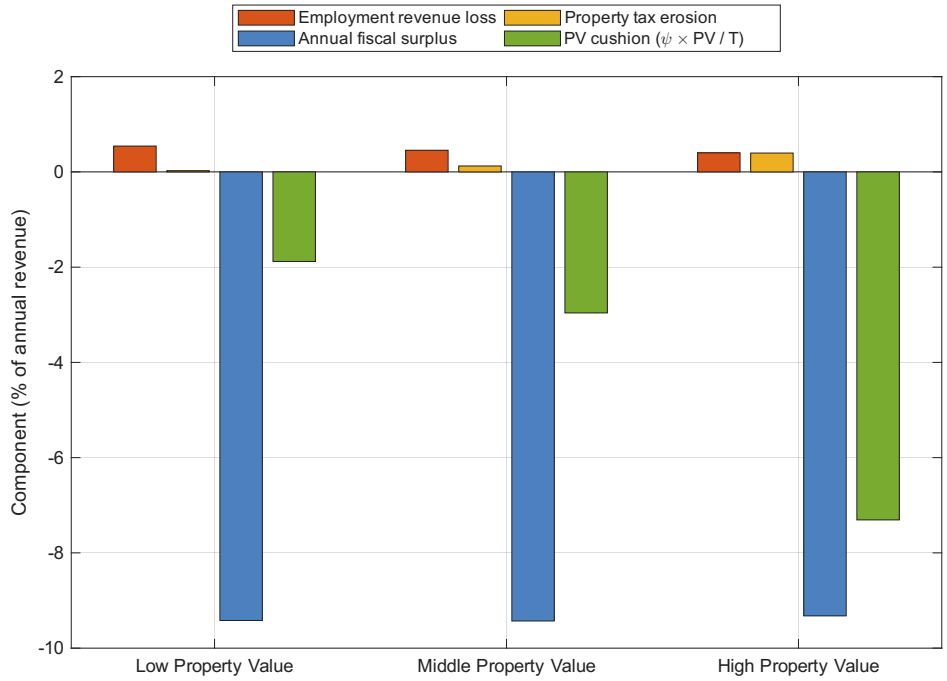


FIGURE IV
Annual Fiscal Loss Decomposition by Property Value Tercile

Notes: Decomposition of the annual fiscal loss into four components: employment revenue loss (red), property tax erosion (orange), annual fiscal surplus (blue, shown as negative = protection), and annualized PV emergency cushion (green, $\psi \times \text{PV}_c/\overline{\text{PV}}/T$). Values are percent of annual revenue, conditional on a negative shock realization. The fiscal surplus is the dominant buffer ($\sim$9%) and is similar across terciles. The PV cushion is the marginal differentiator: high-property-value cities receive nearly 4$\times$ the annualized emergency capacity of low-property-value cities. Calibrated at $\psi = 0.380$, $T = 9.4$ yr.

A calibrated model of stochastic industrial dynamics and fiscal default rationalizes these estimates. The model generates the observed asymmetry between high- and low-property-value cities through three reinforcing channels: shock diversification (concentrated cities have higher employment volatility), property tax stability (high-value cities are less sensitive to sector contractions), and latent fiscal capacity (wealthy cities can tap property wealth as an emergency buffer). With only two free parameters, the model matches the key regression moments closely.

Several caveats warrant mention. First, the 20-sector NAICS classification may understate true concentration if substantial heterogeneity exists within supersectors; finer industry classifications could reveal concentration that the supersector aggregation masks. Second, the model treats fiscal ratios as exogenous, abstracting from the possibility that cities optimally

adjust tax and expenditure policies in response to employment shocks; endogenizing this margin is a natural extension. Third, the analysis focuses on general-purpose municipal issuers and may not generalize to special-purpose entities (e.g., school districts, utilities) that face different revenue structures. Fourth, I address alternative explanations including liquidity, management quality, and institutional constraints, the paper does not model the endogenous fiscal policy responses—such as reserve accumulation or revenue diversification—that cities may adopt in response to employment concentration risk.

These findings have implications for both investors and policymakers. For municipal bond investors, employment concentration—conditional on property values—provides information about credit risk beyond what is captured by standard fiscal metrics and credit ratings. For local policymakers, the results suggest that economic development strategies aimed at diversifying the employment base may yield fiscal benefits through lower borrowing costs, particularly for cities with limited property wealth. The interaction between diversification and property values implies that one-size-fits-all prescriptions are inadequate: the marginal value of diversification is highest precisely where fiscal capacity is thinnest.

References

Manuel Adelino, Igor Cunha, and Miguel A. Ferreira. The economic effects of public financing: Evidence from municipal bond ratings recalibration. *Review of Financial Studies*, 30 (9):3223–3268, 2017.

Kenneth R. Ahern. The business of city hall. Working Paper, University of Southern California, 2024.

Kenneth R. Ahern, Lei Kong, and Xinyan Yan. The anatomy of concentration: New evidence from a unified framework. Working paper, University of Southern California, 2026.

Andrew Ang, Vineer Bhansali, and Yuhang Xing. The muni bond spread: Credit, liquidity, and tax. Working paper, Columbia University, 2014.

Joshua D. Angrist and Jörn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press, 2009.

Cristina Arellano. Default risk and income fluctuations in emerging economies. *American Economic Review*, 98(3):690–712, 2008.

David H. Autor, David Dorn, and Gordon H. Hanson. The china syndrome: Local labor market effects of import competition in the United States. *American Economic Review*, 103(6):2121–2168, 2013.

Jonathan Babcock, Jason G. Cummins, and Christopher A. Parsons. Is *HBCU* debt mispriced? *Working Paper*, 2019.

Tania Babina, Chotibhak Jotikasthira, Christian T. Lundblad, and Tarun Ramadorai. Heterogeneous taxes and limited risk sharing: Evidence from municipal bonds. *Review of Financial Studies*, 34(1):509–568, 2021.

Timothy J. Bartik. Who benefits from state and local economic development policies? Book, W.E. Upjohn Institute for Employment Research, Kalamazoo, MI, 1991.

Olivier Jean Blanchard and Lawrence F. Katz. Regional evolutions. *Brookings Papers on Economic Activity*, 1992(1):1–75, 1992.

Nicholas Bloom and John Van Reenen. Measuring and explaining management practices across firms and countries. *Quarterly Journal of Economics*, 122(4):1351–1408, 2007.

Kirill Borusyak, Peter Hull, and Xavier Jaravel. Quasi-experimental shift-share research designs. *Review of Economic Studies*, 89(1):181–213, 2022.

Alexander W. Butler, Larry Fauver, and Sandra Mortal. Corruption, political connections, and municipal finance. *Review of Financial Studies*, 22(7):2673–2705, 2009.

Alexander W. Butler, Larry Fauver, and Sandra Mortal. Aging and public financing costs: Evidence from U.S. municipal bond markets. *Journal of Public Economics*, 214:104731, 2022.

John Capeci. Local fiscal policies, default risk, and municipal borrowing costs. *Journal of Public Economics*, 53(1):73–89, 1994.

Sudheer Chava, Baridhi Malakar, and Manpreet Singh. Impact of corporate subsidies on borrowing costs of local governments: Evidence from municipal bonds. *Review of Finance*, 28(1):117–161, 2024.

Kimberly Cornaggia, Jess Cornaggia, and Ryan D. Israelsen. Credit ratings and the cost of municipal financing. *Review of Financial Studies*, 31(6):2038–2079, 2018.

Joshua D. Coval and Tobias J. Moskowitz. Home bias at home: Local equity preference in domestic portfolios. *Journal of Finance*, 54(6):2045–2073, 1999.

Gerrit de Wit. Firm size distributions: An overview of steady-state distributions resulting from firm dynamics models. *International Journal of Industrial Organization*, 23(5–6): 423–450, 2005.

Darrell Duffie and Kenneth J. Singleton. Modeling term structures of defaultable bonds. *Review of Financial Studies*, 12(4):687–720, 1999.

Dennis Epple and Chester Spatt. State restrictions on local debt: Their role in preventing default. *Journal of Public Economics*, 29(2):199–221, 1986.

Xavier Gabaix. Zipf's law for cities: An explanation. *Quarterly Journal of Economics*, 114 (3):739–767, 1999.

Pengjie Gao, Chang Lee, and Dermot Murphy. Municipal borrowing costs and state policies for distressed municipalities. *Journal of Financial Economics*, 132(2):404–426, 2019.

Pengjie Gao, Chang Lee, and Dermot Murphy. Financing dies in darkness? The impact of newspaper closures on public finance. *Journal of Financial Economics*, 135(2):445–467, 2020.

Robert Gibrat. *Les Inégalités Économiques*. Librairie du Recueil Sirey, Paris, 1931.

Paul Goldsmith-Pinkham, Isaac Sorkin, and Henry Swift. Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624, 2020.

Paul Goldsmith-Pinkham, Matthew T. Gustafson, Ryan C. Lewis, and Michael Schwert. Sea level rise exposure and municipal bond yields. Working Paper 30660, NBER, 2022.

Government Finance Officers Association. Fund balance guidelines for the general fund. `https://www.gfoa.org/materials/fund-balance-guidelines-for-the-general-fund`, 2015. Best Practice, originally issued 2002, revised 2015.

Richard C. Green, Burton Hollifield, and Norman Schürhoff. Financial intermediation and the costs of trading in an opaque market. *Review of Financial Studies*, 20(2):275–314, 2007.

Adam M. Guren, Alisdair McKay, Emi Nakamura, and Jón Steinsson. Housing wealth effects: The long view. *Review of Economic Studies*, 88(2):669–707, 2021.

Lawrence E. Harris and Michael S. Piwowar. Secondary trading costs in the municipal bond market. *Journal of Finance*, 61(3):1361–1397, 2006.

D. Roderick Kiewiet and Kristin Szakaly. Constitutional limitations on borrowing: An analysis of state bonded indebtedness. *Journal of Law, Economics, and Organization*, 12 (1):62–97, 1996.

Hayne E. Leland. Corporate debt value, bond covenants, and optimal capital structure. *Journal of Finance*, 49(4):1213–1252, 1994.

David A. Lesmond, Joseph P. Ogden, and Charles A. Trzcinka. A new estimate of transaction costs. *Review of Financial Studies*, 12(5):1113–1141, 1999.

Francis A. Longstaff. Municipal debt and marginal tax rates: Is there a tax premium in asset prices? *Journal of Finance*, 66(3):721–751, 2011.

Erzo G. J. Luttmer. Selection, growth, and the size distribution of firms. *Quarterly Journal of Economics*, 122(3):1103–1144, 2007.

Byron Lutz, Raven Molloy, and Hui Shan. The housing crisis and state and local government tax revenue: Five channels. *Regional Science and Urban Economics*, 41(4):306–319, 2011.

Robert C. Merton. On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance*, 29(2):449–470, 1974.

National League of Cities. City fiscal conditions 2025. Technical report, National League of Cities, Washington, DC, 2025.

Robert Novy-Marx and Joshua Rauh. Fiscal imbalances and borrowing costs: Evidence from state investment losses. *American Economic Journal: Economic Policy*, 4(2):182–213, 2012.

Wallace E. Oates. The effects of property taxes and local public spending on property values: An empirical study of tax capitalization and the Tiebout hypothesis. *Journal of Political Economy*, 77(6):957–971, 1969.

James M. Poterba. Tax reform and the market for tax-exempt debt. *Regional Science and Urban Economics*, 19(4):537–562, 1989.

Jennifer Roback. Wages, rents, and the quality of life. *Journal of Political Economy*, 90(6): 1257–1278, 1982.

Sherwin Rosen. Wage-based indexes of urban quality of life. *Current Issues in Urban Economics*, pages 74–104, 1979.

Paul Schultz. The market for new issues of municipal bonds: The roles of transparency and limited access to retail investors. *Journal of Financial Economics*, 106(3):492–512, 2012.

Michael Schwert. Municipal bond liquidity and default risk. *Journal of Finance*, 72(4): 1683–1722, 2017.

Douglas Staiger and James H. Stock. Instrumental variables regression with weak instruments. *Econometrica*, 65(3):557–586, 1997.

James H. Stock and Motohiro Yogo. Testing for weak instruments in linear IV regression. *Identification and Inference for Econometric Models*, pages 80–108, 2005.

Charles M. Tiebout. A pure theory of local expenditures. *Journal of Political Economy*, 64 (5):416–424, 1956.

U.S. Census Bureau. Annual survey of state and local government finances. *Census of Governments*, 2017.

Internet Appendix

Appendix A. Cross-Sectional Evidence

Because employment concentration is primarily a cross-sectional characteristic, a natural first test examines whether the HHI–spread relationship holds in a pure between-city regression. I collapse the panel to a single cross-section by averaging all variables to city-level means over the sample period (2005–2022), then estimate OLS regressions of mean spreads on mean HHI with the same control variables used in the panel regressions: log median property value, general obligation bond share, callable share, average years to maturity, state income tax rate, unemployment rate, log trades, log population, and log par value traded. State fixed effects absorb time-invariant state-level tax and institutional differences. Standard errors are clustered by state (41 clusters).

Table IA.1 reports the results for 1,165 cities. Column (1) confirms that cities with higher average employment concentration have higher average bond spreads, even after controlling for observable city characteristics and state fixed effects. Columns (2)–(3) interact HHI with property value terciles formed from city-mean median property value; to avoid collinearity, log median property value is excluded from the controls in these specifications. The cross-sectional gradient is monotonically decreasing: the HHI coefficient for low-property-value cities is approximately 96 basis points with state fixed effects, compared to 10 basis points for high-property-value cities.

These cross-sectional results confirm that the between-city variation in HHI is associated with meaningful spread differences, particularly among fiscally vulnerable cities. The panel regressions in the main text add city fixed effects that exploit only within-city variation over time, providing a more demanding identification strategy that controls for all time-invariant city characteristics.

Appendix B. Robustness: Predetermined Property Value Terciles

A potential concern with the contemporaneous property value tercile interaction is reverse causality: a negative employment shock could simultaneously raise HHI, depress property values, and increase spreads, mechanically sorting cities into lower terciles. To address this concern, Table IA.2 replicates the main OLS and IV specifications using property value terciles formed from the *prior year's* median property value. Because the lagged tercile assignment is predetermined relative to the current year's outcomes, it is unaffected by contemporaneous shocks. The results are qualitatively similar to the main estimates: the

TABLE IA.1
Cross-Sectional OLS: City-Mean Spreads on City-Mean HHI

	(1) Baseline	(2) HP × HHI	(3) HP × HHI
HHI	59.8		
	(45.8)		
HHI × Low Property Value		149.8***	96.0*
		(49.8)	(51.8)
HHI × Medium Property Value		108.3**	77.9*
		(47.0)	(44.2)
HHI × High Property Value		76.2	10.4
		(48.3)	(36.6)
Log median home value	-14.3***		
	(4.4)		
GO bond share	-14.0**	-4.3	-16.2**
	(6.1)	(4.6)	(6.0)
Callable share	13.3	10.3	12.4
	(9.3)	(12.9)	(9.2)
Avg years to maturity	2.2***	2.9***	2.3***
	(0.8)	(0.7)	(0.8)
State income tax rate	3.3	2.0***	3.8*
	(2.1)	(0.4)	(2.1)
Unemployment rate	3.9***	6.7***	5.0***
	(1.3)	(1.0)	(1.2)
Log trades	15.7***	17.0***	15.6***
	(3.2)	(3.9)	(3.1)
Log population	4.3*	6.9**	4.9*
	(2.4)	(2.9)	(2.7)
Log par value	-20.0***	-22.5***	-20.4***
	(3.7)	(3.8)	(3.7)
State FE	Yes	No	Yes
SE clustering	State	Robust	State
Cities	1,163	1,165	1,163
R^2	0.381	0.284	0.375

Notes: Dependent variable: city-mean maturity-matched yield spread (bps). Each observation is a city, with variables averaged over all panel years (2005–2022). Columns (2)–(3) interact HHI with property value terciles formed from city-mean median home value; log median home value is excluded from controls to avoid collinearity. State fixed effects absorb state-level tax and institutional differences. Standard errors clustered by state (41 clusters) in columns with state FE; heteroskedasticity-robust otherwise. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

HHI–spread gradient is concentrated among low-property-value cities, with attenuated effects for high-property-value cities. The lagged high-property-value interaction is −96 basis points (SE 69), compared to −149 in the contemporaneous specification—a 35% attenuation that suggests a portion of the contemporaneous heterogeneity reflects endogenous sorting of shocked cities into lower terciles. The qualitative pattern—a steep gradient for low-property-value cities and a moderated gradient for high-property-value cities—is preserved.[11]

TABLE IA.2
Robustness: Predetermined (Lagged) Property Value Terciles

	OLS	2SLS
HHI	112.3*	483.8**
	(54.4)	(208.3)
HHI × Medium HP (lagged)	-81.0**	-104.7**
	(37.8)	(46.8)
HHI × High HP (lagged)	-96.0	-118.3
	(68.8)	(134.5)
Observations	17511	16282
Fixed effects	City, Year	City, Year
First-stage F		9.68

Notes: Dependent variable: maturity-matched yield spread (bps). All specifications include city and year fixed effects. Property value terciles are formed within each year's cross-section using the PRIOR YEAR's median home value, addressing potential endogeneity of contemporaneous tercile assignment. Column (2) instruments HHI and HHI×tercile interactions with Bartik predicted values. Standard errors two-way clustered by city and year. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

APPENDIX C. ROBUSTNESS: STATE × YEAR FIXED EFFECTS

A potential concern is that time-varying state-level factors—such as changes in state fiscal policy, state bond market conditions, or state-level exposure to national economic trends—could confound the HHI–spread relationship. The baseline specifications include separate city and year fixed effects, which absorb time-invariant city characteristics and common annual shocks, but do not control for state-level factors that vary over time. State × year fixed effects absorb all such time-varying state-level factors, providing a more demanding test that identifies the HHI–spread relationship solely from within-state, within-year variation across cities.

[11]A fully time-invariant tercile (fixed at the earliest available year) is absorbed by city fixed effects, eliminating the identifying variation in the interaction term. The lagged tercile provides the appropriate middle ground: predetermined relative to contemporaneous shocks while retaining sufficient time variation for identification.

Table IA.3 reports OLS estimates with city and state × year fixed effects (replacing the separate city and year fixed effects used in the main tables). This more demanding specification identifies the HHI–spread relationship solely from within-state, within-year variation across cities. Column (1) reports the baseline HHI coefficient. Column (2) interacts Industrial Concentration with annual property value terciles, paralleling the main specification in Table III.

Appendix D. Robustness: Additional Liquidity Controls

A potential concern is that employment concentration proxies for secondary-market illiquidity rather than credit risk: concentrated cities tend to be smaller, with fewer outstanding bonds and thinner trading, which could inflate spreads through a liquidity premium rather than a default premium. The baseline specifications control for log number of trades and log par value traded, but these are coarse measures of market liquidity.

To address this concern more rigorously, I construct four additional liquidity measures from the trade-level MSRB data:

(1) *Bid-ask proxy* (Harris and Piwowar, 2006): for each bond on each trading day where both dealer-to-customer sales and dealer purchases from customers occur, I compute the mean yield differential. The city-year measure is the median across all qualifying bond-days.
(2) *Yield dispersion* (Green et al., 2007): the standard deviation of yields across all trades of the same bond within a year, averaged across bonds within each city-year. Higher dispersion indicates greater price uncertainty and information asymmetry.
(3) *Zero-trading-day fraction* (Lesmond et al., 1999): for each bond, the fraction of trading days in a year with no recorded trades, averaged across bonds within each city-year. Higher values indicate more infrequent trading.
(4) *Turnover ratio*: total par value traded divided by total par amount outstanding (from EMMA security characteristics), aggregated to the city-year level. Higher turnover indicates more active secondary markets.

Table IA.4 reports the results in two panels. Panel A reproduces the baseline specification (corresponding to Column 1 of Table III) and sequentially adds each liquidity measure. Panel B reproduces the property value interaction specification (corresponding to Column 2 of Table III) and adds the same liquidity controls. In each panel, Columns (2)–(5) add one liquidity measure at a time and Column (6) includes all four simultaneously.

TABLE IA.3
Robustness: State × Year Fixed Effects

	Baseline	HP Interaction
Industrial Concentration	63.3	170.7***
	(39.8)	(42.4)
Ind. Concentration × Medium Property Value		-153.6***
		(37.3)
Ind. Concentration × High Property Value		-186.5***
		(56.9)
Medium Property Value (tercile)		24.6***
		(5.4)
High Property Value (tercile)		30.8***
		(7.8)
Log Median Property Value	-12.4	-14.5
	(9.2)	(9.6)
GO Bond Share	-32.9***	-32.6***
	(5.5)	(5.5)
Callable Share	8.2	8.4
	(12.2)	(12.1)
Average Years to Maturity	2.6***	2.6***
	(0.4)	(0.4)
State Income Tax Rate	—	—
Unemployment Rate	2.1**	2.1**
	(0.9)	(0.9)
Log Number of Trades	6.3**	6.3**
	(2.6)	(2.6)
Log Population	-15.4	-14.7
	(11.1)	(11.0)
Log Par Value Traded	-9.2***	-9.2***
	(2.2)	(2.2)
Observations	18823	18823
Fixed effects	City, State×Year	City, State×Year
Within R^2	0.052	0.054

Notes: The dependent variable is the maturity-matched bond yield spread in basis points. All specifications include city and state × year fixed effects, absorbing all time-varying state-level factors. Column (1) reports the baseline HHI coefficient. Column (2) interacts Industrial Concentration with annual property value terciles (base category = lowest third). Standard errors (in parentheses) are two-way clustered by city and year. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The HHI coefficient is remarkably stable across all specifications in both panels. In Panel A, the point estimate ranges from 51 to 87 basis points and, if anything, *increases* when liquidity controls are added—rising from 54 in the baseline to 87 in the kitchen-sink specification. This pattern implies that the omitted liquidity variables were creating a modest downward bias on the HHI coefficient: concentrated cities happen to have some liquidity characteristics (e.g., lower yield dispersion) that reduce spreads, partially masking the true credit-risk effect. Controlling for liquidity reveals more of the HHI–spread gradient, not less. Panel B confirms that the property value interaction pattern is equally robust: the concentration premium remains concentrated among low-property-value cities regardless of which liquidity controls are included.

Among the liquidity measures, yield dispersion and the bid-ask proxy are the strongest predictors of spreads, and the within-R^2 increases from 0.070 to 0.083 when all four liquidity controls are included. Crucially, this additional explanatory power does not come at the expense of the HHI coefficient: the concentration–spread relationship is not a liquidity artifact.

Appendix E. IV Diagnostic Tables

This appendix reports supporting tables for the two-channel IV design in Section IV.

V.A. Flow Validation

Table IA.5 reports first-difference regressions validating the channel separation of the two instruments. The key requirement for identification in the two-channel design is that each instrument generates variation in its target endogenous variable without contaminating the other channel. To test this, I regress annual changes in each endogenous variable on each instrument separately, with the same controls and fixed effects as the main specification. Column (1) confirms that the annual Guren shock strongly predicts $\Delta \ln(\text{PV})$ (coefficient 0.135, $t = 8.91$), while column (2) shows it has zero predictive power for ΔHHI (coefficient 0.000, $t = 0.04$). Symmetrically, column (3) confirms that ΔBartik predicted HHI predicts ΔHHI (coefficient 0.054, $t = 2.39$), and column (4) shows it has no correlation with $\Delta \ln(\text{PV})$ (coefficient -0.000, $t = -0.00$). The channel separation is clean in first differences, validating the instrument construction.

TABLE IA.4
Liquidity Robustness: HHI Coefficient with Additional Liquidity Controls

	(1) Baseline	(2) +Bid-ask	(3) +Yield disp	(4) +Zero-trade	(5) +Turnover	(6) All liq
Panel A: Baseline						
Industrial Concentration	54.0	84.2**	52.0	51.3	52.5	87.2**
	(39.8)	(42.0)	(40.6)	(40.1)	(39.8)	(41.1)
Bid-ask proxy		4.8				11.2***
		(3.3)				(3.3)
Yield dispersion			10.9**			23.2***
			(4.3)			(5.0)
Zero-trading-day fraction				-240.8**		-281.7**
				(114.0)		(121.7)
Turnover ratio					14.8	5.1
					(13.6)	(11.7)
Observations	18,858	16,741	18,557	18,858	18,855	16,713
Within R^2	0.070	0.075	0.072	0.072	0.071	0.083
Panel B: Property Value Interaction						
Industrial Concentration (Low PV)	143.6**	172.5***	146.4***	143.0***	142.0***	177.7***
	(52.8)	(56.2)	(53.9)	(53.9)	(52.4)	(56.2)
Ind. Conc. × Medium PV	-133.3***	-125.5**	-141.5***	-138.6***	-133.3***	-126.1**
	(49.1)	(54.2)	(48.9)	(51.9)	(48.9)	(55.9)
Ind. Conc. × High PV	-148.6**	-158.0*	-151.6**	-149.8**	-148.2**	-164.7**
	(62.4)	(80.9)	(64.3)	(64.1)	(62.2)	(81.9)
Observations	18,858	16,741	18,557	18,858	18,855	16,713
Within R^2	0.072	0.077	0.074	0.074	0.073	0.085

Notes: The dependent variable is the maturity-matched bond yield spread in basis points. All specifications include city and year fixed effects and the full set of controls from Table I. Standard errors (in parentheses) are two-way clustered by city and year. Panel A reports the baseline HHI coefficient. Panel B reports the HHI coefficient interacted with annual property value terciles (base = lowest third). Columns (2)–(5) each add one liquidity measure; column (6) includes all four simultaneously. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

V.B. Weak-Instrument Diagnostics

Table IA.6 reports identification tests and weak-IV-robust inference for the preferred specification. Panel A tests the joint system with two endogenous variables. The Kleibergen-Paap rk Wald F-statistic is 13.7, well above the Stock and Yogo (2005) 10% critical value of 7.03, rejecting the null of weak instruments. The underidentification test (KP rk LM = 7.96, $p = 0.005$) confirms that the instruments are relevant. The Anderson-Rubin test, which is

TABLE IA.5
Flow Validation: Annual Changes

	$\Delta \ln(\text{PV})$ (1)	ΔHHI (2)	ΔHHI (3)	$\Delta \ln(\text{PV})$ (4)
Annual Guren HP Shock	0.135***	0.000		
	(0.015)	(0.003)		
Δ Bartik Predicted HHI			0.054**	−0.000
			(0.023)	(0.064)
Unemployment Rate	−0.013***	−0.000	−0.000	−0.013***
	(0.001)	(0.000)	(0.000)	(0.001)
log(Population)	0.084***	0.001	0.001	0.086***
	(0.014)	(0.001)	(0.001)	(0.014)
State Income Tax Rate (%)	0.006***	0.000**	0.000***	0.008***
	(0.001)	(0.000)	(0.000)	(0.001)
Within R^2	0.063	0.000	0.002	0.057
Fixed effects	City, Year	City, Year	City, Year	City, Year
Observations	15,183	15,183	15,183	15,183

Notes: This table validates the instruments by regressing annual changes on the excluded instruments. Column (1) confirms that the Guren house price shock predicts $\Delta \ln(\text{PV})$ (relevance). Column (2) shows that the Guren shock does not predict ΔHHI (exclusion). Column (3) confirms that ΔBartik predicted HHI predicts ΔHHI (relevance). Column (4) shows that ΔBartik does not predict $\Delta \ln(\text{PV})$ (exclusion). Standard errors clustered by city are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

robust to weak instruments (Staiger and Stock, 1997), rejects the null that both endogenous variables have zero coefficients ($F = 3.72$, $p = 0.046$). Panel B reports single-channel diagnostics: the Bartik instrument has a first-stage F of 12.0 for HHI, and the Guren instrument has a first-stage F of 8.7 for ln(PV), indicating adequate strength in each channel individually.

TABLE IA.6
Weak-Instrument Diagnostics

	Statistic	p-value
Panel A: Joint system (2 endogenous variables)		
Kleibergen-Paap rk Wald F	13.7	
Stock-Yogo 10% critical value	7.03	
Kleibergen-Paap rk LM (underidentification)	7.96	0.0048
Anderson-Rubin F (joint $\beta_1 = \beta_2 = 0$)	3.72	0.0458
Panel B: Single-channel IV		
KP F: Bartik $\rightarrow$ HHI	12.0	
AR F: HHI channel ($\beta_{\text{HHI}} = 0$)	1.43	0.2483
KP F: Guren $\rightarrow$ ln(PV)	8.7	
AR F: Property value channel ($\beta_{\text{PV}} = 0$)	3.97	0.0625
Fixed effects		City, Year
Clustering		City, Year
Observations		16,367

Notes: Controls include unemployment rate, log(population), and state income tax rate. Panel A tests the joint system with 2 endogenous variables and 2 excluded instruments. Panel B instruments each endogenous variable separately, controlling for the other. Anderson-Rubin tests are robust to weak instruments (Staiger and Stock, 1997). Standard errors two-way clustered by city and year.