The Geography of the Local Service Multiplier

Rachel Bryant[1]

University of Chicago

November 15, 2026

Please click here for the latest version.

Abstract.
State and local governments spend over $20 billion a year on firm-specific subsidies, rationalized by the direct jobs these firms create and the indirect jobs their payroll supports. This paper studies, conditional on firm entry, the strength of the local service multiplier and how it depends on commuting and shopping networks in a city. Using employer-employee linked Census data, I find that about half of the entrant's new employees were not previously attached to the city labor market. Whether these workers migrate, commute from elsewhere, or were previously unemployed determines how much of the entrant's payroll becomes new local spending. I build a spatial equilibrium model across neighborhoods, calibrated to these rates of migration, commuting, and labor force participation, in which residents' choices of where to shop determine where their income is spent. In the model, the service multiplier is smaller than most existing estimates. For firm entries of the same relative size, the multiplier is about 2.2 times as large at the 90th percentile as at the 10th. Aggregate characteristics of the local economy explain only a small share of this variation, and much of the rest reflects the commuting and shopping flows around the entrant's specific location. Finally, I combine the model with a new dataset of subsidized firm entry events that records promised jobs, realized jobs, and the state and local portions of a subset of subsidy awards. In most events, the added residential property tax from the entry covers the local share of the award, and residents' welfare gains exceed the award's full cost, assuming the firm would not have entered otherwise. However, the same subsidy can deliver substantially different welfare gains at other locations in the same city.

[1]I am grateful to my advisors Milena Almagro, Erik Hurst, and Esteban Rossi-Hansberg for their guidance, support, and encouragement. This paper benefited from comments from participants in the UEA North America Meeting, UEA Summer School at LSE and the Trade & Spatial Working Group at the University of Chicago. Any views expressed are those of the author and not those of the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 3321 (CBDRB-FY27-P3321-R13471). This research uses data from the Census Bureau's Longitudinal Employer-Household Dynamics Program, which was partially supported by the following National Science Foundation Grants SES-9978093, SES-0339191 and ITR-0427889; National Institute on Aging Grant AG018854; and grants from the Alfred P. Sloan Foundation. All errors are my own.

1 Introduction

Over 40 percent of U.S. workers are employed in service industries,[2] whose demand is primarily supported by the local population. When a positive economic shock raises wages and employment for residents, part of their additional income is spent on services, making service employment reflective of local economic conditions more so than tradable employment. Tradable firms depend on the local economy in terms of factor prices but less for product demand because they often sell to a wider regional market. As a result, tradable firms are the primary beneficiaries of economic development subsidies from state and local governments that total between $20-46 billion a year, the largest category of place-based spending in the United States. Policymakers often cite this local spending channel as a spillover of attracting tradable firms:

> "A lot of them [new workers] live within the city of Reed City. So, they own homes here and pay property taxes. But even if they're not here inside the city, they're doing business here. They're supporting our grocery stores, they're supporting our gas stations. They're supporting our local restaurants."
>
> —Mayor of Reed City, Michigan

The response of service employment, which I call the service multiplier, matters for assessing subsidized tradable firm entry because a subsidy's expected benefit is the probability it changes a firm's location decision multiplied by the gain to residents. A larger service multiplier raises that gain, so the subsidy can leave residents better off even when it is less likely to sway the firm to locate there. A large empirical literature, discussed below, studies how local economies respond to new tradable employment, from broad shocks such as industry-wide growth and federal programs to individual entrants such as million-dollar manufacturing plants, foreign multinationals, and high-skilled firms. Most of these studies measure effects within a fixed area, such as cities and counties, so less is known about what determines the scale and location of the service response. It is natural to think of the service multiplier as an interaction between the entry shock and its location that jointly determine how easily the firm hires new workers and how much of the new income translates into local service spending. Therefore, in this paper, I ask: (1) How do commuting and shopping patterns shape the spatial incidence of the local service multiplier and its resulting size? (2) How does the service multiplier change the cost-effectiveness of firm-specific subsidies?

To answer these research questions, I build and estimate a quantitative spatial model in the spirit of Miyauchi et al. (2025) that distinguishes commuting flows from shopping flows between Census tracts within commuting zones (CZs) in the U.S. The granularity of the model is essential because commuting and shopping are two networks that operate on different spatial scales: individuals are

[2] 43.3 percent of employed civilians ages 25 and older in the 2024 American Community Survey 1-Year Public Use Microdata Sample, including part-time workers, public employees, and the self-employed. Service sector refers to NAICS 44–45, 61, 62, 72, 81. Over half of these workers are in education and health care (NAICS 61, 62), where demand follows the number of nearby residents but payment comes largely from governments and insurers

nearly twice as likely to shop within a ten-minute drive of home as to work there.[3] Additionally, the firms I examine are among the largest half a percent of establishments in the United States, yet the median subsidized firm employs less than one percent of its county's workers and approximately 0.16 percent of its commuting zone's labor force. Because of the scale mismatch, the impact of firm entry is evident in tract-level employment and residential maps, but it is obscured by county- and CZ-level economic trends.

I start by assembling data to translate an entrant's payroll into locations of increased service spending. I first use the Census Bureau's Longitudinal Business Database (LBD) to identify large single-establishment tradable entries and expansions and where each one locates. Then I use the Longitudinal Employer-Household Dynamics (LEHD) to link each entrant to the workers it hires. I have their earnings and where they live before and after being hired, so I can distinguish residents from commuters and migrants and measure how much of the payroll is new local income. I use cell phone ping data, which record trips from each home tract to service establishments, to examine the shopping flows that determine where that income is spent. Finally, I build a new dataset of subsidized firm entries that records each firm's site, its promised and realized jobs, and, for a subset of awards, the state and local portions of the subsidy. Together, these datasets allow me to follow an entrant's payroll from the firm to its workers' homes and to the places they shop at the level of the Census tract, and to compare the resulting gains with local costs.

Using these data, I provide stylized facts about who tradable entrants hire and where those workers spend, and new reduced-form evidence on the response of service employment. Following each entrant's hires through their employment and residence histories observed in the LEHD, I find that the jobs an entrant creates are a poor approximation for the income increase of residents. Only 50 percent of hires previously worked in the CZ, while 28 percent live outside it, 14 percent move in, and 8 percent were unemployed or self-employed. Each group brings a different amount of new income. A migrant brings a full paycheck, a commuter only what is spent near work, and a worker hired away from a local employer only a raise unless that employer refills the job or raises pay to keep its other workers. Among incumbent employers, the tradable industries most exposed to the entrant's hiring lose a modest amount of employment, while nontradable industries show no clear change.

About 48 percent of after-housing consumer spending goes to locally consumed services, and cell phone ping data show that individuals start 74 percent of their shopping trips from home and 26 percent from work. A third of all trips end in an individual's home tract or an adjacent tract. Another third of trips end at fixed shopping hubs within the CZ, and a third go to neither. The local third is the part that varies across tracts, because residents throughout the CZ shop at the same hubs. I therefore use the local part to measure a tract's exposure to an entrant. I assign each hire's service spending to the tract where that individual lived before entry and to the tracts that

[3] Among 2019 trips within a commuting zone, 42 percent of service visits and 23 percent of jobs are within a ten-minute drive of the home census tract. Visits are from cell phone ping data (Veraset) for retail, education, accommodation and food, and other services (NAICS 44, 45, 61, 72, 81). Jobs are LODES commuting flows. Drive times are from OSRM.

border it, so a tract is more exposed when more of the entrant's hires live in it or next to it. In a stacked difference-in-differences design that ranks tracts by this exposure, in CZs with no other large entrant, the most exposed quartile of tracts gains 32 service jobs and $1.9 million in service payroll relative to the least exposed quartile seven years after entry. Gains of this size in a small number of tracts are difficult to detect in county totals, yet they are evidence that even these large, tradable firm entry events can alter the distribution of service employment.

In the model, service and tradable firms compete for a common pool of workers and commercial floor space. Congestion in the labor market is weakened by the availability of commuters, migrants, and nonemployed locals. However, these three groups differ in what share of their earnings represents new income and thus new service demand for a city. Monte et al. (2018) show that a productivity shock raises employment in places more open to commuting, but the effect on service employment is less clear. Places that are more open in terms of commuting flows are also more open in terms of shopping flows (the correlation between the two is 0.78). Because the majority of service spending is close to home, new commuters bring less new service demand than migrants or unemployed residents. Meanwhile, migrants bring increased demand for housing, which raises costs and can reduce service spending.

I model firm entry as a productivity increase in the entrant's tract and sector, with the increase scaled so that employment rises by the firm's promised number of jobs. Individuals travel to consume services, incurring a cost that makes them sensitive to distance. Tradable firms at other locations shrink as the entrant increases equilibrium wages and commercial rents, consistent with the modest employment decline in the tradable industries most exposed to the entrant's hiring I observe in the data. Service firms are monopolistically competitive and enter freely. Both residents and firms pay rents for land in separate markets, and residents earn rental income through ownership of residential land while commercial rents accrue to absentee landlords. The model yields welfare, the fiscal cost of a job, and the share of the response that stays inside the jurisdiction that pays for it. The model isolates the effect of the entering firm, removing confounding trends that obscure the empirical comparisons at the county and CZ level.

I quantify the model separately for each of the firm entrant events in its CZ. The migration and sector-choice elasticities are estimated with shift-share designs of industry growth instrumenting for cross-CZ variation in earnings, and the residence-workplace pair choice and shopping elasticities are estimated from gravity equations on tract-pair flows. I calibrate the inbound CZ commuting and labor participation elasticities by minimum distance to the hiring shares I observe in the employer-employee linked data. The remaining spending, cost, and supply parameters come from the national accounts, the shopping trip data, and the literature. Given these parameters, I invert the model to recover amenities, productivities, and fixed costs so that it reproduces each commuting zone's economy before entry. None of these steps directly uses the difference in service employment across tracts, so the event-study estimates are untargeted moments. To validate the model, I solve the equilibrium with and without each entrant, run the difference-in-differences design on the model's output, and compare the model-implied coefficients with their empirical counterparts. The model's

predictions fall within one standard error of the estimated effect on service jobs and within one and a half standard errors of the estimated effect on service payroll.

With the model quantified, I return to the two questions, taking into account the full general equilibrium effects. First, how do commuting and shopping shape the size of the service multiplier and where it lands? For every 100 jobs an entrant promises, the median commuting zone gains 29 service jobs, or 58 per 100 net tradable jobs as measured in the literature. It gains 80 jobs in total once the 49 jobs lost at other tradable firms are counted. Residents' welfare rises by about $11.9 a year for every resident of the CZ, or about $7.9 million in total. Of each $100 of additional payroll at the entrant's site, residents' earnings rise by about $41, and about $31 reaches service firms in the CZ. Migrants bring the most new spending per worker, while commuters and residents who enter the labor force bring less but ease the competition for workers. The shopping network, and in particular workplace origin shopping, is important for the spatial incidence of the multiplier. More than half of the entrant's workers live outside the firm's county and they shop mostly near home. As a result, a majority of the county's service revenue gain is from workplace origin shopping trips, which make up only a quarter of all shopping trips in the data.

There is significant variation in the service multiplier across firm entry shocks, ranging from about 18 to 39 service jobs per 100 entrant jobs between the 10th and 90th percentiles. Market characteristics explain only about 13 percent of the variation. Larger CZs, as measured by baseline employment, have higher wages, so each additional promised job brings more payroll. However, these higher wages are also paid to service employees, so each dollar of service revenue supports fewer jobs. Larger markets also have less elastic housing, so the firm entry shock attracts fewer new residents. These effects offset and there is not a statistically significant relationship between the market size and the size of the service multiplier for shocks of the same job size. Where the firm locates within the county matters more. Within the typical county, the site at the 90th percentile adds about 1.7 times as many service jobs as the site at the 10th. A firm near the county border hires more workers who live in the neighboring county, and those workers spend most of their service budget near home, so both their earnings and their spending leave the host county. A firm in the middle of the county hires a larger fraction of its new workers from county residents, so the county loses less than half as many service jobs to its neighbors as it would if the firm located near the border.

Finally, how does the service multiplier change the cost-effectiveness of a subsidy? It lowers the cost of a job, but the county that pays keeps only part of the gain, because a county is more open to commuting and shopping than the commuting zone around it. The award data I build separate the local share of each subsidy from the state share for about one fourth of entry events. At the median, the county subsidy forgoes about $2,900 a year per net CZ job, against about $4,700 if the service jobs are left out. The county holds 28 percent of the CZ's residents and receives 61 percent of the gain in resident income and 62 percent of the service job gain. Median resident welfare per public dollar is therefore $3.1 across the CZ but $1.1 in the host county, which bears the entire cost. When the firm would have entered regardless, the county is simply forfeiting tax revenue. I find

that the subsidy benefits CZ residents only if it causes at least 12 percent of the entries it subsidizes, a threshold well within the 2 to 25 percent that Bartik (2018) reports for typical incentives.

Contribution to the literature. This paper makes three contributions. First, it measures the local multiplier of individual subsidized plants and traces it from payroll to service employment. Estimates from industry-wide growth (Moretti, 2010; van Dijk, 2017), from shift-share designs (Bartik, 1991; Bartik and Sotherland, 2019; Autor et al., 2013, 2021; Hershbein and Stuart, 2024; Yagan, 2019; Hornbeck and Moretti, 2024; Acemoglu and Restrepo, 2020), and from regional variation in military procurement (Nakamura and Steinsson, 2014) treat the county, the CZ, or the state as a single point and do not observe who fills the new jobs (Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022). Plant-opening and subsidized-firm designs (Greenstone et al., 2010; Patrick, 2016; Patrick and Partridge, 2022; Slattery and Zidar, 2020; Hanson and Rohlin, 2024) observe the firm but not its workers, and county employment net of the entrant is too noisy to recover the effect of a single firm whose hiring is a small fraction of total employment. Evaluations of larger place-based programs do find local gains. Manufacturing employment in Tennessee Valley Authority counties was still growing faster than in comparison counties four decades after federal transfers lapsed (Kline and Moretti, 2014b), and federal Empowerment Zone designation raised employment in zone neighborhoods by 12 to 21 percent (Busso et al., 2013), although California's enterprise zones did not raise employment (Neumark and Kolko, 2010). These programs fund infrastructure, tax credits, or block grants across a region or neighborhood rather than a single firm. I follow the entrant's hires to where they live and shop and add up the links in equilibrium. The results reconcile the two literatures. An entrant's payroll supports local service jobs, but only about half of its own jobs are net additions, because it crowds out other tradable producers. The service response is spread over the neighborhoods where its workers live and shop, and 38 percent of it lands outside the host county, rendering a county-wide comparison incomplete. Because much of the payroll goes to workers who already had jobs in the CZ or who live outside it, the response to a single plant depends on the new income its payroll creates rather than on its headcount.

Second, I show that how open a local market is depends on where its boundary is drawn, and I measure the multiplier at the boundary of the county that pays as well as at the boundary of the CZ. The model has the neighborhood structure of Ahlfeldt et al. (2015) and Tsivanidis (2026) and adds two margins that the entry question requires. Individuals shop from home and from work, as in the consumption-trip model of Miyauchi et al. (2025) and in line with smartphone evidence that consumption trips are short and concentrated (Davis et al., 2019; Agarwal et al., 2020; Oh and Seo, 2023). I add a sector choice where individuals switch between tradable and service employment, so that the entrant's wage pressure can directly impact service firms. Both factor demand and output demand operate in a CZ whose boundary workers and shoppers cross in both directions, with leakage measured from data. Monte et al. (2018) show that local employment elasticities depend on commuting openness. I add shopping openness, which differs from commuting openness at every boundary, and show that the county that pays for an award is considerably more open than

the CZ that contains it. The model is solved for a single large employer entering fine geography in 95 actual CZs rather than across counties.

Third, I separate the part of the multiplier that depends on the firm type and the part that depends on the local characteristics of the market, namely the commuting and shopping openness and the housing supply elasticity. I measure each channel by comparing the CZ total to the county that finances the employment. Theory treats mobility as a continuous relocation margin (Kline and Moretti, 2014a; Busso et al., 2013; Fajgelbaum and Gaubert, 2020), and the incidence of state business taxes is divided among workers, landowners, and firm owners with aggregate elasticities (Suárez Serrato and Zidar, 2016, 2023). Garin (2025) argues that local job counts are uninformative when the jobs go to migrants, commuters, or workers who would have been employed anyway, and the contrast between location-based and hiring-based incentives in enterprise zones (Neumark and Kolko, 2010; Freedman, 2013) shows that the concern is practical. I contribute to this analysis by showing that while award programs measure performance by the jobs a firm promises, most of the variation in what those jobs contribute locally depends on the commuting and shopping openness of the place as well as housing flexibility.

The rest of this paper is organized as follows. Section 2 describes the policy setting and the data. Section 3 presents three facts that follow the entrant's payroll to whom it hires, where those workers shop, and where service employment responds. Section 4 develops the model and Section 5 quantifies it. Section 6 measures the net employment and welfare effects of entry, shows what shopping, commuting, and migration add, traces where the gains land within the commuting zone, and asks how much the multiplier varies across markets and across locations within a county. Section 7 asks what each job costs the county that subsidizes it. Section 8 concludes.

2 Data

Firm entry. In the motivating empirics, I examine all large tradable entries and expansions. Using the Longitudinal Employer-Household Dynamics (LEHD) data, which link 95% of all private-sector workers to their employer, earnings, and residence, I identify large openings and expansions of tradable establishments in the 29 U.S. states for which I have data. I trace the hiring of each firm to the neighborhoods where its workers live, and then to the neighborhoods where they shop. I measure the response of nontradable employment, payroll, and establishments in those locations. Subsidized deals enter this sample when they meet the same size and industry screen I use to classify large, tradable firms in the public data.

The LEHD sample contains approximately 400 such entrants. I classify an entrant into high- and low-skill-intensive tradables.[4] Table 1 compares the two firm types. I find that high-skill-intensive tradable firms locate in substantially larger markets, with a median population of 1.21 million, compared with 291,000 for low-skill-intensive entrants. High-skill-intensive firms locate closer to the CZ center (defined as the group of tracts with the highest employment totals), with

[4] A firm is classified as high skill when its four-digit industry had an above-median share of workers with a bachelor's degree in 2000 and as low skill otherwise.

a median driving time of 18.2 minutes to the downtown area rather than 25.5 minutes. Lastly, median annual pay indicates that the entry of a high-skill-intensive firm delivers a much larger payroll shock, but its entry is less common in less densely populated labor markets.

Table 1: Characteristics of low- and high-skill entrants

	Low-skill entrants	High-skill entrants
Firms (disclosure-rounded)	200	200
Hires		
Annual pay of high-skill hires (2023 $)	57,420	130,900
Annual pay of low-skill hires (2023 $)	44,690	87,240
Site and labor market		
Driving time to CZ center (minutes)	25.54	18.17
CZ population	290,600	1,211,000
Annual CZ population growth, 3 years before entry	0.68	0.61

Notes: Released estimates from the LEHD entrant sample. Characteristics are medians across entrants of each type; entrant counts are disclosure-rounded. Pay refers to the entrant's hires by skill and is annual, in 2023 dollars. Proximity to the CZ center is measured in driving minutes. Population growth is the annual percentage change in the CZ from five to two years before entry.

Subsidized firm entry. I assemble awards from Good Jobs First, state development agencies, and county records. I match each subsidized firm to its establishment record in Data Axle, a commercial dataset of U.S. businesses, and verify the location of the subsidized site with the Google Maps API. This matching allows me to ensure the firm actually enters. I then observe the exact address of its entry and can assign it an industry code.[5] For approximately 40% of the awards, I also separate the portion offered locally, as property taxes forgone by the county and other local governments, from the portion paid by the state through credits, rebates, and grants. For the remaining awards, I simply retain the amount spent by the state. I measure the jobs a firm actually creates from state compliance reports and other documented employment records. After location site verification, the subsidy dataset contains 216 firm-site events (Appendix A.1). Among the 64 firm entry events with reported employment at least four years after award, 39 percent of firms reach their promised employment, the median firm delivers 79 percent of the jobs promised, and 8 firms more than double their hiring promise (Figure A.1).

I focus on the entry and expansion of large, tradable firms. Consistent with the subsidy literature, I consider a firm "large" if it promises at least 500 new jobs, or 0.50% of the total employment in its commuting zone (CZ). I distinguish tradable industries, whose demand originates primarily in national or international markets, from nontradable industries, which mainly serve local demand (Moretti, 2010; van Dijk, 2017). Throughout the paper, the service sector whose response I measure

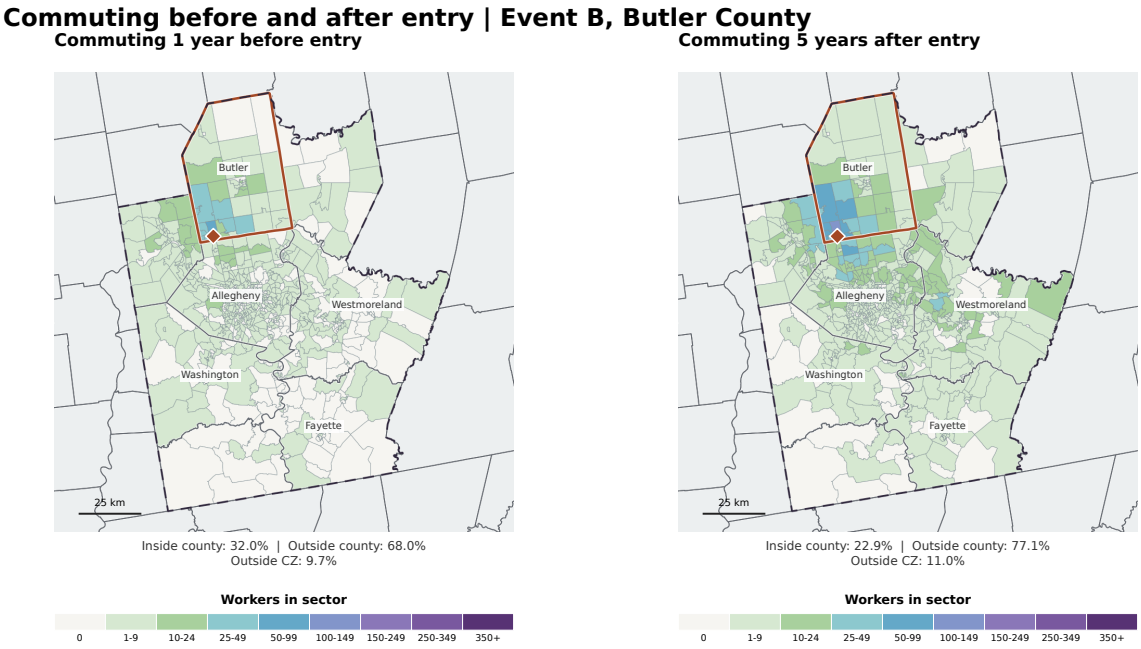
[5]Good Jobs First is a nonprofit research center in Washington, DC, that tracks economic development subsidies. Its Subsidy Tracker database compiles company-specific awards from federal, state, and local programs using government disclosures and public records requests. Data Axle, formerly Infogroup, maintains a commercial establishment-level database of U.S. businesses compiled from public sources and verified by telephone. Its annual snapshots, available from 1997, record each establishment's name, location, industry, employment, sales, year established, and parent company.

comprises retail, education, health, food services, and other services (NAICS 44–45, 61, 62, 722, and 81).[6]

Local markets and jurisdictional boundaries. I use the commuting zone (CZ) to define the local labor market, and set Census tracts as the unit for locations within it. I use CZs because they are contiguous county groupings (unlike MSAs) that maximize residence–workplace connections within them (Tolbert and Sizer, 1996), but neither a county nor a CZ is a closed economy. Across entrants, the mean share of workers who live in the same county is 46%, and at the CZ boundary it rises to 78%.

Figure 1 shows both boundaries for one entry into a suburban county of a large CZ that shows a more extreme split between county and CZ. It maps where the people working in the entrant's tract live, one year before the firm enters and five years after. The entrant locates at the edge of its county, next to the CZ's central county. 77 percent of the added workers live outside the host county, but only 11 percent of the tract's workers live outside the CZ five years after entry. Section 6 discusses this entry using the model.

Figure 1: Where the workers at an entrant's tract live, before and after entry



Notes: Public LODES origin–destination counts of primary jobs (JT01) in the broad All Other Services group (SI03) at the entrant's tract, by the worker's residence tract, one year before and five years after entry. These counts include every employer in that group in the tract. This firm entered a business park with the entry of other employers within successive years. The diamond marks the entrant's tract and the outlined area its host county. The denominator for the shares is all workers, so the Outside CZ is included in the Outside County share.

[6]Because LODES and the Quarterly Workforce Indicators report only two-digit sectors, the public-data outcomes and the model's baseline use all of NAICS 72, which adds accommodation; the national spending shares in Section 3 also include it.

Shopping trips. To measure where residents shop, I use 2019 smartphone visit records from Veraset (Veraset, 2022). I weight each device equally within a month, so that the annual measure is a device-weighted distribution of visits from each home tract to each destination tract. Shopping is local: 84.80% of service visits stay within the shopper's home commuting zone, yet only 10.60 percent of within-CZ visits stay in the home tract (Table 3 in Section 3 and Appendix Table A.12), reflecting the need for a model with shopping trips to capture effects at granular scales. Because the visit data are observed in 2019, after all of my firm entry events, I use them to describe where entry-linked demand may be observable and to estimate shopping gravity. They are not used in the inversion of the model.

Prices, employment, and wages. Land data come from CoreLogic deed and assessor records over 2004–2020, which give residential and commercial repeat-sales and hedonic price indexes, building square footage, and property counts for each tract. Public employment outcomes come from LODES, which records jobs by workplace tract, resident workers by home tract, and the commuting flows between them. I use the Quarterly Workforce Indicators for county earnings by education that I allocate to tracts as workplace wages, and I use the County Business Patterns to provide establishment counts by county and industry. Appendix A documents each source and its construction.

3 Motivating Facts

Fact 1. Only part of entrant payroll becomes additional local income. I use LEHD records for approximately 400 large tradable firm entries and expansions observed between 2004 and 2017, and I classify the new workers by their residence. Table 2 shows where these hires live and what they earn. Together these determine how much of new payroll becomes new local income.

Combining the hiring counts and earnings summarized in Table 2 gives an approximate account of where entrant payroll ends up. Of every $100 in new payroll, only about $32 is new income for the CZ's residents. About $28 goes to workers who live outside the CZ, and another $40 replaces what local hires earned at their previous jobs. Of the new $32, $15 is paid to workers who move into the CZ and $17 is the net earnings gain of hires who already lived there.[7]

The $40 that replaces hires' previous pay is not necessarily lost to the CZ. When a hire's former employer refills the vacated job with someone who was not working in the CZ or increases wages to retain employees, that pay is new income too, so the $32 of new income is a lower bound. To examine what happens to existing firms, I compare incumbent industries in the entrant's commuting zone with the same industries in five matched commuting zones. An industry's exposure is the entrant's employment multiplied by the share of national job-to-job movers who flow between that three-digit

[7]The earnings allocation is an approximation. It weights the joint residence–workplace individual counts by their median earnings, pooling prior industries and individual skills within entrant type and prior workplace. Then I scale the prior and after-hire allocations to the earnings totals within entrant type and post-hire residence. I assume hires with no observed prior job have zero prior income.

Table 2: Entrant hires, earnings, and payroll

A. Residence

	Percent of hires
Continuing CZ residents	57.79
Live outside the CZ	27.90
Migrated to the CZ	14.32
	Percent of continuing residents
No observed prior job	9.56

B. Earnings (2023 $ thousands)

	After hire	Increase
Continuing CZ residents	81.8	20.1
Live outside the CZ	84.2	26.6
Migrated to the CZ	87.1	27.8
No observed prior job	60.5	—

C. Payroll by residence one year after hire

	Percent of entrant payroll
Paid to residents inside the CZ	71.69
Paid to residents outside the CZ	28.31

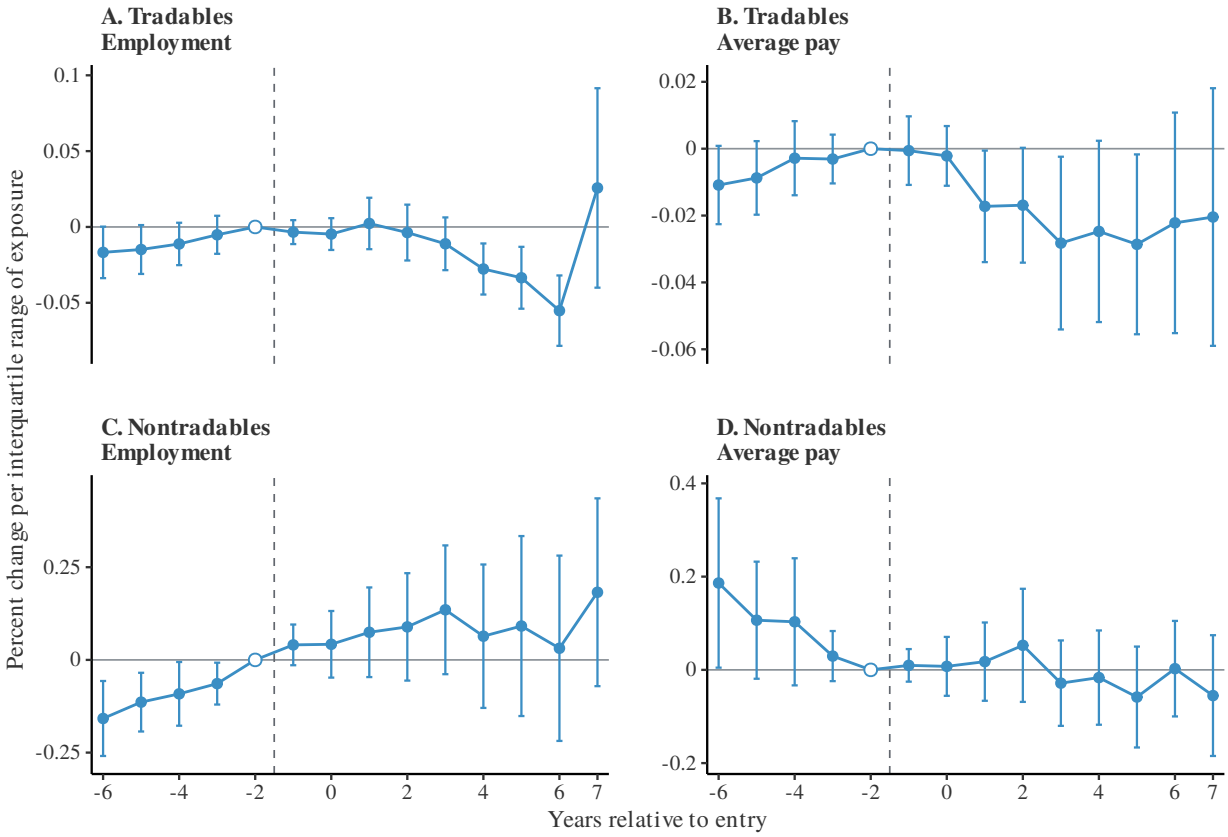
Notes: Released LEHD tabulations pooling all firm entries and expansions. Panel A reports shares of hires by residence one year after hire. The no-observed-prior-job share uses continuing residents as its denominator. Panel B reports worker-weighted averages of the earnings medians, in thousands of 2023 dollars, for the Panel A groups. The outside row excludes the 1 percent of hires who lived in the CZ before hire and outside it after. After-hire earnings are annual earnings one year after hire; the increase is the difference between the after-hire and prior-year median earnings. Panel C uses group earnings totals by post-hire residence. Shares may not sum to 100 because of rounding.

industry and the entrant's industry before entry, divided by the industry's baseline employment in the CZ. I estimate an event study of log employment and log average pay on this exposure interacted with event time and an indicator for the entrant's CZ, with event–CZ–industry, event–CZ–year, and event–industry–year fixed effects, so each coefficient compares more and less exposed industries in the entrant's CZ relative to the same industries in the matched CZs (Appendix A.4). Tradable industries more exposed to the entrant lose employment four to six years after entry, by up to 0.06 percent per interquartile range of exposure, while nontradable industries show no clear change (Figure 2). Appendix A.4 shows the raw employment and pay trajectories of the establishments directly exposed to the entrant.

Fact 2. Local income is spent on services close to home. Fact 1 separates entrant payroll from new income for the CZ's residents. The next question is how much of that income reaches service firms in the CZ, and where.

Workers will travel across a commuting zone for their job, but they are less willing to travel for groceries or a haircut. To directly compare commuting and shopping, I estimate, separately for

Figure 2: Incumbent industries linked to the entrant through national worker flows



Notes: Predicted labor exposure uses national job-to-job flows between each incumbent three-digit industry and the entrant's industry, averaged over the three years before entry. Outcomes are log employment and log average pay for tradable (top row) and nontradable (bottom row) incumbent industries. Released log coefficients and 95 percent confidence intervals are multiplied by 100, so they read as approximate percent changes for one interquartile range of exposure. Year -2 is omitted, the dashed line falls between years -2 and -1, and each panel has its own scale. The release covers 80 CZs.

each commuting zone c and each travel mode $m \in \{S, C\}$, a gravity equation for shopping visits (S) and commuting flows (C) between tracts i and j,

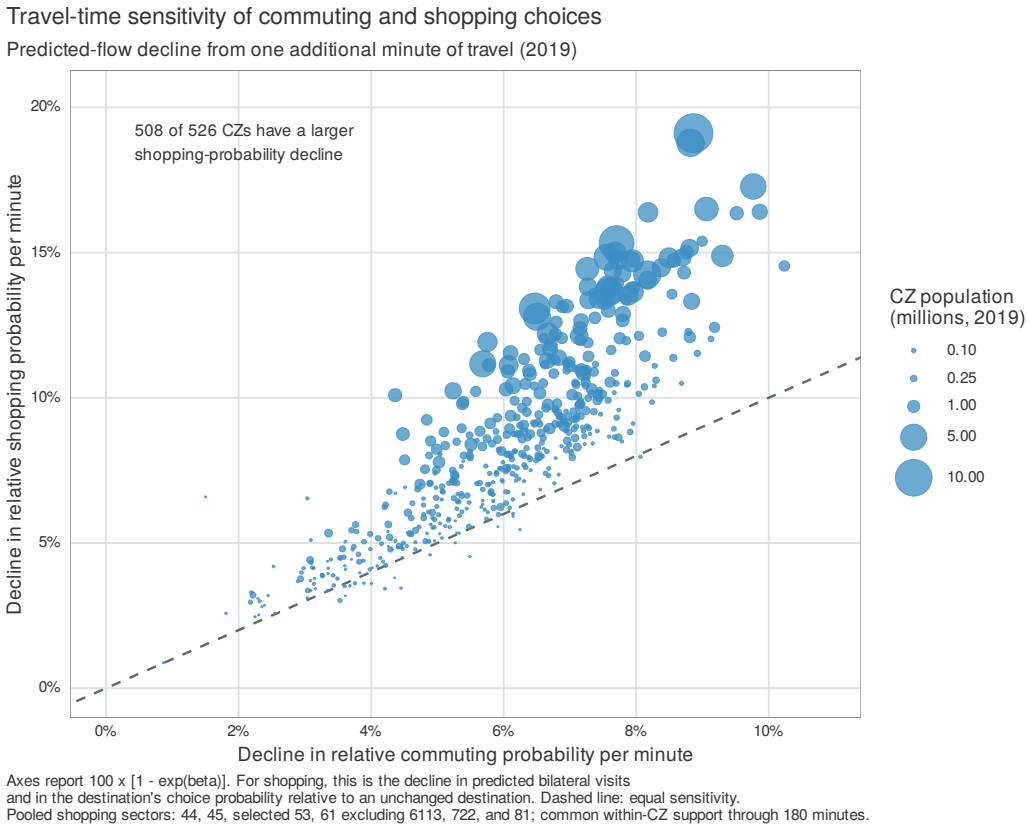
$$\mathbb{E}\left[F_{ijc}^{m}\right] = \exp\left(\alpha_{ic}^{m} + \delta_{jc}^{m} + \beta_{c}^{m}\, t_{ijc}\right), \qquad m \in \{S, C\}, \tag{1}$$

where F_{ijc}^{S} is 2019 device-weighted service visits from residents of i to establishments in j, F_{ijc}^{C} is the 2019 LODES count of workers living in i and working in j, and t_{ijc} is driving time in minutes. I estimate equation (1) by Poisson pseudo-maximum likelihood (Silva and Tenreyro, 2006). The origin effects α_{ic}^{m} absorb the number of residents and workers in each tract, and the destination effects δ_{jc}^{m} absorb workplace size and destination attractiveness.

Figure 3 plots the two estimates for the 526 commuting zones. Shopping decays faster than commuting in 97% of CZs. At the median, an additional minute of driving time reduces predicted shopping flow by 8.1% and predicted commuting flow by 6.1%, a difference of a third. With an

increase of a ten-minute drive, shopping flows fall by 57% and commuting flows by 47%. The two coefficients are strongly correlated across places (0.86), so places in which commuting is sensitive to travel time also have concentrated shopping. These results show that a new workplace can pull employees from a larger area of the commuting zone while the consumption their earnings generate remains concentrated near a smaller set of residential neighborhoods.

Figure 3: Travel-time sensitivity of shopping and commuting across commuting zones



Notes: Each observation is a commuting zone. The horizontal axis is the percent decline in predicted commuting flow from one additional minute of driving time, and the vertical axis is the analogous decline in service visits, both computed as $100[1-\exp(\widehat{\beta}_c^m)]$ from the CZ-specific PPML estimates of equation (1). The dashed line denotes equal decay. The point size is proportional to 2019 population. Shopping pools NAICS 44–45, 61 excluding colleges, 722, and 81. Both trips use 2019 flows and 2020 OSRM driving times.

Within a CZ, pooling across service sectors, 11 percent of shopping visits end in the shopper's home tract and 24 percent in an adjacent tract, so over a third of service trips end within one tract boundary of a person's home tract. Nearly half of the rest goes to shopping hubs, clusters of destination tracts within ten minutes' drive of one another that each receive at least 2 percent of their CZ's service visits (Appendix A.5). Hubs receive 43.4 percent of within-CZ visits, and two-thirds of hub visits are to a hub outside the shopper's home neighborhood.[8]

[8]The number of hubs rises with market size, with an elasticity with respect to population of about 0.43 under the 2 percent definition (Appendix A.5).

In Fact 3, I compare service employment across tracts within a CZ, which requires shopping exposure that varies with where an entrant's hires live. Because residents throughout a CZ shop at the same hubs, spending routed through the full visit matrix ends in similar tracts irrespective of origin. Among the 85 firm entry events in CZs with more than one entry, destination alone explains 76.6 percent of the variation in that exposure. A design built on the full matrix picks up growth at each market's hubs but cannot attribute it to a particular entrant (Appendix A.6).[9] Instead, I use the local part that does vary. 34.4 percent of a tract's visits end in its own or an adjacent tract, against 3.2 percent if it shopped like its CZ as a whole (Appendix A.5). Fact 3 therefore assigns each worker's spending to their home neighborhood, 30 percent to the home tract (its share of local visits in Table 3) and 70 percent split equally among the adjacent tracts to approximate the shares of the home and adjacent tract splits.

Table 3: Allocation of within-CZ service visits between the home neighborhood and hubs, 2019

Sector	Same tract	Adjacent tract	Nonlocal 2% hub
NAICS 44	12.2	26.1	26.4
NAICS 45	9.0	23.6	33.5
NAICS 722	9.6	23.1	30.5
NAICS 81	13.5	22.9	22.7
Combined core	*10.6*	*23.8*	*29.3*
NAICS 61	7.2	16.2	32.8

Notes: Entries are percentages of annual device-weighted visits, conditional on the destination remaining within the device's home CZ. The categories are mutually exclusive. The remaining visits go to other destinations within the home CZ. "Adjacent tract" excludes the home tract; "Nonlocal 2% hub" is a shopping hub that excludes both the home tract and its adjacent tracts. Sector definitions are in Appendix A.5.

Fact 3. Service employment grows near the residential locations of workers. The first two facts imply that an entrant's payroll may be identifiable in service employment in the neighborhoods where its workers live. I estimate this effect on service employment and payroll with a stacked event study of tracts (Cengiz et al., 2019; Deshpande and Li, 2019). Every tract in a stack is exposed to the same entry, so each stack compares more and less exposed tracts rather than treated and untreated ones. Because the two groups receive equal weight within each stack and every stack receives the same total weight, the pooled coefficient is the simple average of the stack-specific contrasts (Gardner, 2021; Wing et al., 2024).

I define firm entry as the first year in which the firm records employment, and I measure a tract's exposure to this entry from the number of the entrant's hires who lived in it or an adjacent

[9]The visit data are from 2019, after every entry in the sample. If shopping itself responds to entry, because new stores open near the entrant or near its workers' homes, exposure routed through hubs and the full network would partly reflect that response. The home-and-adjacent rule does not have this problem, because which tracts it ranks as most exposed depends only on where hires lived before entry and on which tracts border theirs.

tract before entry, weighted by the 30/70 shares above. Within each event, a tract is treated if its exposure is in the top quartile of the event's tracts and serves as a control if it is in the bottom quartile but still has positive exposure. Mean exposure differs between the two groups by 23 workers. I restrict the sample to 40 entries in commuting zones with no other observed entry, so that no control tract is exposed to a second event.

For each event, I build a stack of its top- and bottom-quartile tracts from six years before entry to seven years after. I estimate, for employment and payroll in retail, food services, and other services (NAICS 44, 45, 722, and 81),[10]

$$Y^o_{ejt} = \alpha^o_{ej} + \delta^o_{et} + \sum_{k \in \{-6,\ldots,7\} \setminus \{-2\}} \theta^o_k D_{ej} \mathbb{1}\{t - \tau_e = k\} + u^o_{ejt}, \tag{2}$$

where Y^o_{ejt} is service employment (jobs) or annual service payroll (millions of 2023 dollars) in tract j in year t for event e, D_{ej} indicates a top-quartile tract, and τ_e is the entry year. α^o_{ej} and δ^o_{et} are tract and calendar-year fixed effects, so a tract is compared only with tracts in the same market in the same year. The reference year is two years before entry rather than one, because hiring can begin in the year before the entrant first records employment.[11] Within each event the two groups receive equal weight, and standard errors are clustered by CZ, because tracts in the same market share local shocks.

Under parallel trends and no anticipation, θ^o_k is the effect of entry on service employment in highly exposed tracts relative to less exposed tracts in the same market, k years after entry (Callaway et al., forthcoming). Reading it as the effect of greater exposure further requires that the most exposed tracts would have responded to low exposure as the least exposed tracts did. Exposure is fixed by where hires lived before being hired, so it cannot reflect moves after hiring or later changes in shopping. The stack-specific year effects remove any shock common to the market, so only differences within the market identify θ^o_k. Additionally, the two groups move together before entry. Every pre-entry coefficient for employment and payroll is small, and every 95 percent confidence interval includes zero (Figure 4), which would be unlikely if service firms responded to the announcement before hiring began. The remaining concern is that entrants hire from neighborhoods whose service sectors were about to grow for other reasons. Appendix A.6 also measures exposure by new income and by residences predicted from public commuting data, and gives the full construction of the stacks.

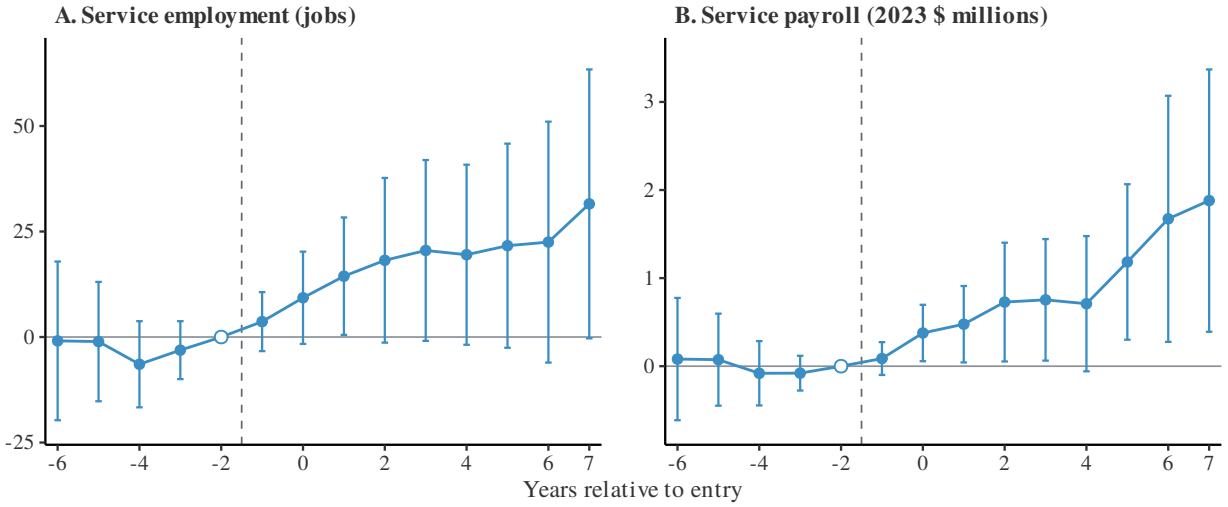
Seven years after entry, a top-quartile tract has $1.9 million more in annual service payroll, in 2023 dollars, than a bottom-quartile tract in the same market, with a 95 percent confidence interval of $0.39 to $3.4 million (Figure 4). Service employment follows the same path and reaches 32 more jobs by year seven, but it is estimated less precisely, with an interval from just below zero to 63

[10] Education and health care (NAICS 61 and 62) are excluded. Their location is tied to schools and hospitals and their funding comes largely from governments and insurers, so their employment responds less directly to nearby household spending.

[11] Employment is measured on March 12, so hiring that begins after March 12 of the preceding year can occur while recorded employment for that year is still zero.

jobs.

Figure 4: Service employment and payroll around entrant workers' pre-entry residences

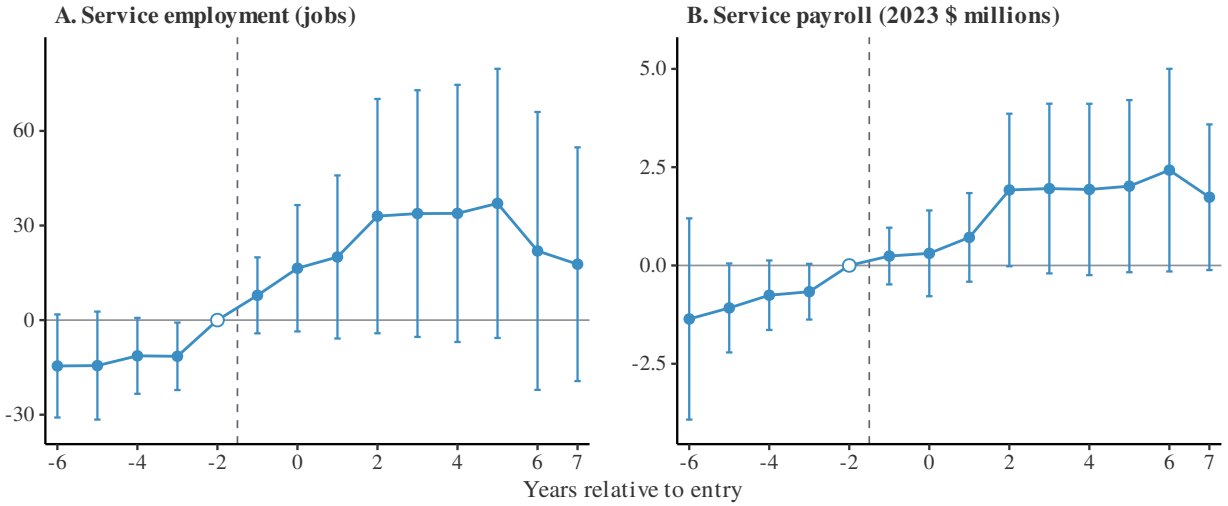


Notes: Stacked event-study coefficients from equation (2), comparing the top and bottom quartiles of pre-entry-resident worker exposure in commuting zones with no other observed events. The difference in mean exposure is 23.5 workers. Each worker's pre-entry home tract receives 30 percent of exposure, and adjacent tracts share 70 percent equally. Outcomes cover retail, food services, and other services (NAICS 44, 45, 722, and 81). Bars are 95 percent confidence intervals; year −2 is normalized to zero. Payroll is displayed in millions of 2023 dollars. The release reports 40 commuting zones.

Figure 5 estimates the same relationship with continuous exposure across all events and exposure tracts, including CZs with more than one observed entry. Exposure from each entry enters as a distributed lag, with endpoint terms that carry exposure from firm entries outside the window, which accommodates several entries of different sizes in the same tract (Schmidheiny and Siegloch, 2023). Distributed lags of this kind are a common way to study policy changes of different sizes (Suárez Serrato and Zidar, 2016; Fuest et al., 2018). This specification uses every entry and the full exposure distribution rather than only those in single-entry CZs, but it assumes that the effect of an entry is proportional to a tract's exposure and that effects add up across entries (Schmidheiny and Siegloch, 2023). Each hire is again placed at their residence before entry, with the same 30/70 split, and the regression includes tract and CZ-by-year fixed effects. The coefficients average 26 service jobs and $2.1 million in annual service payroll per 100 exposed workers over years five through seven. These are slopes across the full exposure distribution, so their size is not directly comparable to the quartile contrast.

From the facts to the model. The three facts follow entrant payroll to workers' homes, to the shops near those homes, and to the service jobs that grow there. They cannot add up to a net effect for the market. Comparisons across neighborhoods mix new jobs with jobs that move from elsewhere in the CZ, and they miss the wages, rents, and prices through which the entrant also crowds out other firms. The model in Section 4 adds those channels. It is calibrated to the hiring

Figure 5: Continuous service responses to pre-entry-residence exposure with 30/70 home-and-adjacent weights



Notes: Continuous event-study coefficients per 100 exposed workers from equation (A.12). Hires are located at their residences before entry. Each home tract receives 30 percent of exposure, and adjacent tracts share 70 percent equally. All events and exposure tracts enter, including CZs with multiple observed events. Outcomes cover retail, food services, and other services (NAICS 44, 45, 722, and 81). The regression includes tract and CZ-by-year fixed effects, is unweighted, and excludes the entrant's workplace tract. Bars are CZ-clustered 95 percent confidence intervals; year −2 is normalized to zero. Payroll is in millions of 2023 dollars. Each panel has its own vertical scale. The release reports 100 commuting zones.

shares in Fact 1 and the shopping gravity in Fact 2, and Section 5 compares its predictions with Fact 3.

4 A Model of Labor and Shopping

I develop a spatial general equilibrium model to measure the service multiplier across neighborhoods and compare it with resident welfare. The economy consists of a single commuting zone (CZ) partitioned into census tracts $i, j \in \mathcal{J}$ differentiated by exogenous amenities and productivity. Individuals choose where to live, work, and shop. Their shopping trips originate from both their home and workplace. The commuting zone is an open economy with workers living just outside its boundary and migration from elsewhere in the country. There are four sectors. Two tradable sectors, high and low skill intensive, produce with Cobb–Douglas technologies and sell at a national price, and two service sectors, also high and low skill intensive, are monopolistically competitive with free entry. Each county levies a property tax to pay for public services, which it buys from outside the CZ, so the revenue leaves the modeled economy. Residents own residential land, while commercial rents accrue to absentee landlords. Residential and commercial floor space are supplied with separate elasticities.

I model an entering tradable firm as a positive productivity shock in a sector and location. This shock increases demand for labor and raises wages, encouraging individuals to join the labor force,

migrate into, or work in the area. A portion of the increase in worker income is spent on local services. Service firms may see revenue gains that translate into increased demand for workers, but in an environment where labor and floor space may be more expensive. The spatial incidence of these net effects is determined by the interaction of workers' choices of where to live, work, and shop, so the gains from job creation accrue not at the firm's entry location but along the commuting and shopping networks that connect it to the rest of the commuting zone.

4.1 Individuals

Individuals differ by skill $k \in \{H, L\}$ and choose where to live, in which sector to work, whether to work, and their residence–workplace pair. They compare living in the CZ, in the accessible outside area, or elsewhere in the country. The order of the region and sector choices differs by skill. High-skill individuals choose their sector first and their region second, so they move between regions more readily than they switch sectors. Low-skill individuals choose their region first and their sector second, so they switch sectors more readily than they move. Within each sector and region, individuals decide whether to work and, if they do, choose a residence–workplace pair. The migration elasticity μ_k differs by skill, while the sector-choice and participation elasticities, ν and η, are common to both.

This nested structure allows the model to distinguish new workers from new consumers. After a positive labor demand shock, substantial shares of new workers migrate, cross the CZ boundary for work, or come from nonemployment (Fact 1). Unemployed residents already consume in the CZ, and those who commute into the CZ spend much of their income outside of it. A semi-open city model with a single migration elasticity could approximate the employment openness of the local market, but it would treat every new worker as a new consumer, and in doing so, overstate the increase in service demand. In this model, these nested employment and location choices determine the labor available to local employers and the locations from which workers originate their shopping trips.

4.1.1 Preferences

Let ℓ index a joint residence and employment choice. For a working individual ω of skill k, this specifies residence i, workplace j, and sector g.[12] The individual consumes housing h, services n, and a tradable good x whose price is normalized to one. For an individual living and working within the CZ, preferences are Cobb–Douglas,

$$U_\ell(\omega) = z_\ell(\omega) \frac{B^R_{ik} B^W_{jgk}}{d_{ij}} h^{\alpha_H} \left(n^{s_N} x^{1-s_N}\right)^{1-\alpha_H}, \tag{3}$$

where B^R_{ik} and B^W_{jgk} are residential and workplace amenities, with workplace attractiveness allowed to differ by sector. Additionally, $d_{ij} = \exp(\kappa t_{ij})$ is the utility cost of a commute of t_{ij} minutes, and

[12] For nonworkers, ℓ specifies residence i and a sector with amenities $B^R_{ik} B^U_{ik}$ and no commute. Appendix B.1 gives the corresponding amenities for choices involving the outside area.

$z_\ell(\omega)$ is an idiosyncratic taste for choice ℓ. The housing expenditure share is α_H, and s_N is the share of remaining income spent on services.[13]

Let y_ℓ denote disposable income, q_i^R the residential rent, and $\mathcal{P}_\ell$ the service price index. Maximizing utility over consumption gives indirect utility $z_\ell(\omega)\Phi_\ell$,[14] where

$$\Phi_\ell = \frac{B_{ik}^R B_{jgk}^W}{d_{ij}} \frac{y_\ell}{(q_i^R)^{\alpha_H} \mathcal{P}_\ell^{(1-\alpha_H)s_N}}. \tag{4}$$

Each individual chooses the option ℓ with the highest $z_\ell(\omega)\Phi_\ell$. The value Φ_ℓ is common to everyone of skill k who makes choice ℓ, while $z_\ell(\omega)$ is the individual's idiosyncratic taste for it. Tastes are drawn independently across individuals from a nested Fréchet distribution and are held fixed across equilibria. For an individual of skill k,

$$\Pr\{z_\ell(\omega) \le t_\ell \text{ for every } \ell\} = \exp\left\{-G_k\left((t_\ell^{-1})_\ell\right)\right\}, \qquad t_\ell > 0, \tag{5}$$

where G_k incorporates the nesting structure and the preference parameters (Appendix B.2). The shape parameters μ_k, ν, η, ψ, and χ govern migration, sector choice, labor-force participation, the choice between jobs inside and outside the CZ, and the residence–workplace choice, respectively.

4.1.2 Employment and location choice nests

Whether an individual moves to the CZ depends on the utility it offers, and that utility depends on the jobs in each sector they can reach from places they would like to live. A higher wage at one workplace raises not only the appeal of commuting there, but also the value of working in its sector, and the value of living in the CZ.

The value of working in sector g sums over every residence–workplace pair available in that sector, including workplaces outside the CZ, and the value of not working sums over residences,

$$\mathcal{W}_{gk} = \sum_{\ell \in \mathcal{L}_{gk}^W} \Phi_\ell^\chi, \quad \mathcal{U}_{gk} = \sum_{\ell \in \mathcal{L}_{gk}^U} \Phi_\ell^\chi. \tag{6}$$

Working and not working then combine into the value of the sector,

$$\mathcal{S}_{gk} = \overline{B}_{gk}^E \mathcal{W}_{gk}^{\eta/\chi} + \overline{B}_{gk}^U \mathcal{U}_{gk}^{\eta/\chi}, \tag{7}$$

where $\overline{B}_{gk}^E$ and $\overline{B}_{gk}^U$ are preference parameters for employment and nonemployment, and the sector values in turn determine the regional and sector choices in equation (8).

Individuals can also live outside the CZ, either in an accessible outside area or elsewhere in the country. The accessible outside area pools the neighboring CZs into a single location whose

[13] For outside residents, the same preferences and demand expressions use the fixed housing share $\overline{\alpha}_H^o$ in place of α_H.

[14] For nonemployment, the amenity component is $B_{ik}^R B_{ik}^U$, income includes nonemployment income y_k^U, and the service price index is P_i, the price of services purchased from the home origin.

residents can work in the CZ, work outside it, or remain nonemployed. Moving between the outside area and the CZ is governed by the same parameter μ_k as migration from elsewhere.

For a resident of the outside area, the choice between a job inside the CZ and one outside it is a separate nest with shape parameter ψ, which governs how readily residents of neighboring CZs take CZ jobs when wages rise. CZ residents who work outside it instead treat the outside area as one more workplace in their residence–workplace choice, so ψ affects only outside residents. Commutes across the CZ boundary have no travel time in the model. The costs of travel into the CZ are part of the location-specific exogenous amenity term calibrated to the observed inbound and outbound flows (Appendix C.6.1). The terms $\mathcal{O}_L$ and $\mathcal{O}_{gH}$ in (8) collect the value of living in the outside area or elsewhere in the country. It can change with entry because outside residents benefit from access to CZ jobs, while utility elsewhere in the country is fixed.

Regional location choice. For low-skill individuals, the outermost nest is the choice whether to live in the CZ; for high-skill individuals, it is the choice of sector. Let $\mathcal{N}_k$ be the fixed national population of skill k and $\overline{B}^S_{gk}$ the preference parameter for sector g. For low-skill individuals, $\overline{B}^R_L$ is the preference parameter for CZ residence and $\mathcal{O}_L$ the value of living in the outside area or elsewhere in the country. For high-skill individuals, both are specific to the sector, $\overline{B}^R_{gH}$ and $\mathcal{O}_{gH}$, and $\mathcal{V}_{gH} = \mathcal{O}_{gH} + \overline{B}^R_{gH}\mathcal{S}^{\mu_H/\eta}_{gH}$ is the value of sector g across regions. The number of CZ residents of skill k in sector g is

$$N_{gk} = \begin{cases} \mathcal{N}_L \dfrac{\overline{B}^R_L \mathcal{Q}_L^{\mu_L/\nu}}{\mathcal{O}_L + \overline{B}^R_L \mathcal{Q}_L^{\mu_L/\nu}} \dfrac{\overline{B}^S_{gL}\mathcal{S}^{\nu/\eta}_{gL}}{\mathcal{Q}_L}, & k = L, \\ \mathcal{N}_H \dfrac{\overline{B}^S_{gH}\mathcal{V}^{\nu/\mu_H}_{gH}}{\sum_{g'\in\mathcal{G}} \overline{B}^S_{g'H}\mathcal{V}^{\nu/\mu_H}_{g'H}} \dfrac{\overline{B}^R_{gH}\mathcal{S}^{\mu_H/\eta}_{gH}}{\mathcal{V}_{gH}}, & k = H, \end{cases} \tag{8}$$

where $\mathcal{Q}_L = \sum_{g\in\mathcal{G}} \overline{B}^S_{gL}\mathcal{S}^{\nu/\eta}_{gL}$. For low-skill individuals, the first factor is the share who live in the CZ and the second is the share of those residents who choose sector g. For high-skill individuals, the order is reversed, so the first factor is the national share who choose sector g and the second is the share of them who live in the CZ. The CZ's population of skill k sums over sectors,

$$N_k = \sum_{g\in\mathcal{G}} N_{gk}. \tag{9}$$

I assume utility elsewhere in the country is fixed. Because the CZ is small relative to the nation, the individuals it draws from do not change wages or prices in the nation. The migration elasticity μ_k is approximately the elasticity of the CZ's population with respect to the value of living there, within each sector for high-skill individuals.[15] When μ_k is low, few individuals move, and a labor demand shock is absorbed by higher wages, prices and rents.

[15]The exact elasticity is $\mu_k(1 - b^{\text{reg}}_k)$, where b^{reg}_k is the share of the national population of skill k living in the CZ or the accessible outside area, computed within sector for high-skill individuals.

Sector choice. Individuals compare employment opportunities in two tradable sector categories (high and low skill intensive) and two service sector categories (high and low skill intensive), so $\mathcal{G} = \mathcal{G}_T \cup \mathcal{G}_N$ with $\mathcal{G}_T = \{T_H, T_L\}$ and $\mathcal{G}_N = \{N_H, N_L\}$.[16] Equation (8) gives the number choosing each sector. The labor demand shock increases wages in one of the tradable sectors, which draws individuals toward tradable employment, increasing labor competition across sectors for incumbent service firms. The choice of N_{gk} includes the ability to not work. Therefore, the sector choice is best described as an employment opportunity choice.

Labor-force participation. Within sector g, individuals compare working with nonemployment.[17]

The number employed in sector g among CZ residents is

$$E_{gk} = N_{gk} \frac{\overline{B}_{gk}^{E} \mathcal{W}_{gk}^{\eta/\chi}}{\mathcal{S}_{gk}}. \tag{10}$$

Higher wages can draw nonworkers into employment, expanding the workforce and raising the income available for local services.

Residence–workplace location choice. Conditional on living in the CZ and working in sector g, an individual of skill k chooses residence–workplace pair $\ell = (i, j, g)$ from the feasible set $\mathcal{L}_{gk}^{W}$ with probability

$$\pi_{\ell k} = \frac{\Phi_{\ell}^{\chi}}{\sum_{\ell' \in \mathcal{L}_{gk}^{W}} \Phi_{\ell'}^{\chi}}, \quad F_{ijgk} = E_{gk} \pi_{\ell k}. \tag{11}$$

The numerator captures the attractiveness of the pair through residential and workplace amenities, disposable income, residential rent, service prices, and commuting costs. The denominator, $\mathcal{W}_{gk}$, sums these utility terms across all feasible residence–workplace choices in the sector, so the probabilities sum to one. Multiplying the probability by the number of employed CZ residents in sector g, E_{gk}, gives the corresponding commuting flow. Nonworkers choose among residences with the same parameter χ.

4.1.3 Income and spending

A working individual supplies one unit of labor and earns the wage w_{jk}^{g} paid to skill k in sector g at workplace tract j (Section 4.3), while a nonworker receives the externally financed payment y_{k}^{U}.

[16] Workers of both skill types can work in either the high or low intensity sectors. These sectors simply employ different shares of each worker type.

[17] Nesting participation within sector treats each individual as attached to a sector even when not working, so that a nonworker compares not working with the within-sector wage. This follows the labor literature on sector-specific skills and attachment. I am assuming human capital is partly specific to an industry or occupation. This follows the literature in which displaced workers who switch industries suffer larger earnings losses (Neal, 1995), and Galle et al. (2023, Table 6), who estimate a larger shape parameter for the choice between nonparticipation and work than for the choice across sectors, consistent with $\eta \geq \nu$.

Current CZ residents also receive a dividend D_k from ownership of residential land.[18] Disposable income is

$$y_\ell = w^g_{jk} + D_k \quad \text{for a worker,} \quad y_\ell = y^U_k + D_k \quad \text{for a nonworker.} \tag{12}$$

Section 4.4 determines the dividend.[19]

At residential rent q^R_i, Cobb–Douglas preferences imply housing demand h_ℓ and service spending γy_ℓ, with

$$h_\ell = \frac{\alpha_H y_\ell}{q^R_i}, \quad \gamma = s_N(1-\alpha_H). \tag{13}$$

Therefore, an additional dollar of disposable income raises service spending by γ.

Workers allocate a share ρ of service spending to purchases originating at home and the remainder to purchases originating at work, with $n = n_i^\rho n_j^{1-\rho}$. Let P_i denote the service price index at origin i. The service price index for workers is $\mathcal{P}_\ell = P_i^\rho P_j^{1-\rho}$. Meanwhile, nonworkers purchase all services from home, so their price index is $\mathcal{P}_\ell = P_i$.

Let y_{ijgk} denote the disposable income of a worker of skill k who lives in i and works in sector g at j, and let U_{ik} be the number of nonemployed individuals of skill k at residence i, summed across sectors. Total service spending originating in tract i is

$$\mathcal{E}_i = \gamma\Big[\underbrace{\rho \sum_{j,g,k} F_{ijgk} y_{ijgk}}_{\text{Workers' purchases from home}} + \underbrace{(1-\rho)\sum_{j,g,k} F_{jigk} y_{jigk}}_{\text{Workers' purchases from work}} + \underbrace{\sum_k U_{ik}\big(y^U_k + D_k\big)}_{\text{Nonworkers' purchases}}\Big], \tag{14}$$

where γ uses the outside housing share for inbound commuters.

4.1.4 Shopping destinations

From each shopping origin i, individuals purchase a CES bundle of services available across destinations j, with elasticity of substitution $\sigma_N > 1$. Tract j offers m_j differentiated service varieties, with prices summarized by p^N_j. Each variety is supplied by a firm operating under monopolistic competition with free entry.[20]

Travel from origin i to destination j takes t_{ij} minutes and raises the effective price of purchases by the iceberg factor $\exp(\kappa t_{ij})$, using the same travel-cost parameter κ as commuting. Under CES demand, these preferences and travel costs enter through $T_{ij} = \zeta_{ij}\exp\{-(\sigma_N - 1)\kappa t_{ij}\}$, which determines each destination's contribution to the service price index. The cost of one unit of the

[18] In the quantification, housing expenditure includes the rental-equivalent value of owner-occupied housing. This essentially is a rent which residents pay to themselves. The assumption matters because rental income stays in the CZ, so part of it is spent on local services. Because the dividend is shared by all current residents, it spreads the gains from higher rents across the population rather than concentrating them on incumbent owners.

[19] Appendix B.3 gives the corresponding income accounting for outside residents.

[20] The model distinguishes two service sectors, N_H and N_L, with separate prices and entry responses. This section presents shopping and revenue aggregated by tract. Appendix B.6 gives the sector-specific equations and derives these aggregates.

service bundle purchased from origin i is[21]

$$P_i = \left[\sum_{j \in \mathcal{J}} T_{ij} m_j (p_j^N)^{1-\sigma_N} \right]^{1/(1-\sigma_N)}. \tag{15}$$

Firm entry expands the range of available services and, holding prices and travel costs fixed, lowers P_i. The expenditure share received by tract j is

$$\pi_{ij}^I = \frac{T_{ij} m_j (p_j^N)^{1-\sigma_N}}{P_i^{1-\sigma_N}}. \tag{16}$$

Individuals direct more spending toward destinations offering more varieties, lower prices, or better travel access. These purchases generate service revenue of $\sum_{i \in \mathcal{J}} \mathcal{E}_i \pi_{ij}^I$ in tract j.

4.2 Service Firms

A service firm in sector $g \in \mathcal{G}_N$ and tract j produces a differentiated variety under monopolistic competition and free entry.[22] With productivity A_{jg}, a firm combines labor $\mathcal{L}$, commercial floor space H, and intermediate inputs X purchased from outside the CZ to produce output Y, using the same Cobb–Douglas form as the tradable sectors (Section 4.3),

$$Y = A_{jg} \, \mathcal{L}^{\alpha_L^g} H^{\alpha_C^g} X^{\alpha_X^g}, \tag{17}$$

where the sector's expenditure shares on labor, commercial space, and intermediates satisfy $\alpha_L^g + \alpha_C^g + \alpha_X^g = 1$.

Under CES demand, firms charge a constant markup over their marginal cost c_{jg},

$$p_{jg} = \frac{\sigma_N}{\sigma_N - 1} c_{jg}. \tag{18}$$

Appendix B.7 derives unit costs and factor demands.

Each firm also incurs an annual fixed cost f_{jg} to maintain its establishment. The fixed cost uses the same inputs as production, so it is proportional to the unit cost c_{jg}, and the input requirement per firm changes with the number of firms m_{jg} in the tract,

$$f_{jg} = \varphi_{jg} c_{jg} m_{jg}^{1/\epsilon_m - 1}, \tag{19}$$

where φ_{jg} scales the requirement. At the CES markup, each firm keeps a share $1/\sigma_N$ of its revenue

[21] Appendix B.6 describes shopping opportunities outside the CZ and derives the expenditure share spent outside the CZ.

[22] Appendix B.7 gives the full conditions for each service sector and their aggregation to tracts.

after variable costs, so free entry drives profits to zero when

$$\frac{R_{jg}}{\sigma_N} = m_{jg} f_{jg}, \quad m_{jg} = \left(\frac{R_{jg}}{\sigma_N \varphi_{jg} c_{jg}} \right)^{\epsilon_m}. \tag{20}$$

The entry elasticity ϵ_m is the elasticity of the number of establishments with respect to revenue relative to unit cost. When $\epsilon_m = 1$, each firm's requirement is constant and establishments rise in proportion to revenue relative to cost. A value below one makes each firm's requirement rise as firms enter, so establishments respond less than proportionally.

Since revenue covers variable and fixed costs, total labor expenditure is $\alpha_L^g R_{jg}$. Allocating this expenditure across skills and dividing by the corresponding wage gives employment:

$$L_{jk}^g = \frac{\alpha_L^g \lambda_{jgk} R_{jg}}{w_{jk}^g}. \tag{21}$$

4.3 Tradable Firms and Entry

Tradable producers in sector $g \in \mathcal{G}_T$ (high and low skill intensive) sell a homogeneous good on the national market at price p^T, which the CZ takes as given.[23]

Tradable production in each tract and sector is competitive, with productivity A_{jg}. Production is Cobb–Douglas in labor, commercial floor space, and intermediate inputs. With the price of intermediate inputs normalized to one, cost minimization gives the unit cost

$$c_{jg} = \frac{(q_j^C)^{\alpha_C^g}}{A_{jg}} \prod_k \left(w_{jk}^g \right)^{\alpha_L^g \lambda_{jgk}}. \tag{22}$$

Production has constant returns and the output price is set nationally, so profits are zero in every location and sector,

$$c_{jg} = p^T, \tag{23}$$

and output Y_{jg} adjusts until the labor and floor-space markets clear. Labor demand is a constant share of revenue,

$$L_{jk}^g = \frac{\alpha_L^g \lambda_{jgk}\, p^T Y_{jg}}{w_{jk}^g}, \tag{24}$$

and demands for commercial space, H_{jg}^C, and intermediate inputs take the same form (Appendix B.8). Equation (23) means that a tradable producer cannot pass higher costs on to its consumers. Given its productivity, each producer can pay only the wages at which it breaks even, and its employment adjusts to the labor supply those wages attract. When the entrant's productivity gain raises the wages workers can earn elsewhere, incumbent tradable producers cannot match them, so they lose workers and shrink.

[23] The high-skill industry comprises NAICS 51, 52, 54, and 55, and the low-skill industry comprises manufacturing (NAICS 31–33). The labels refer to industries' skill intensity, and individuals of either skill can work in either industry.

Entry. I model the entry of a tradable firm as an increase in the productivity $A_{j_E g_E}$ of its sector g_E in its tract j_E, sized so that employment in this tract and sector rises by J^E, the entrant's promised jobs. Higher productivity raises the wages the location-sector can pay, and its employment rises along the upward-sloping labor supply described in Section 4.1. The model does not distinguish the entrant from incumbent producers in the same tract and sector, so J^E measures the location-sector's net growth. Because the shock raises productivity for the whole location-sector, the higher wage it supports is paid to every worker here, not only to new hires. The model has no separate premium for the entrant's workers.

4.4 Land Market Clearing

Each tract has separate markets for residential and commercial floor space. Space cannot be converted between uses, so each tract has two rents. Owners receive r_i^R and r_j^C, and tenants pay these rents plus the property tax rate $t_{a(i)}$ of the county $a(i)$ that contains the tract,

$$q_i^R = (1 + t_{a(i)}) r_i^R, \quad q_j^C = (1 + t_{a(j)}) r_j^C. \tag{25}$$

Floor space is built by combining land, which is fixed in each tract, with structures, so its supply slopes upward in the rent owners receive. I take this supply to have a constant elasticity,

$$H_i^R = \bar{H}_i^R (r_i^R)^{\epsilon_{R,i}}, \quad H_j^C = \bar{H}_j^C (r_j^C)^{\epsilon_C}, \tag{26}$$

as in quantitative urban models with a construction sector (Ahlfeldt et al., 2015). The elasticities measure how much floor space expands when rents rise, and $\epsilon = 0$ gives the fixed supply of Monte et al. (2018). The shifters $\bar{H}_i^R$ and $\bar{H}_j^C$ capture each tract's land area and the ease of building on it.

Residents spend a share α_H of disposable income on housing (13). Let $Y_i^R = \sum_k \sum_{\ell:\, i_\ell = i} n_{\ell k} y_\ell$ denote the total disposable income of tract i's residents, workers and nonworkers alike. Commercial space is demanded by tradable producers, including producers in the entrant's tract and sector, and by service firms. Equating supply and demand in each market gives

$$r_i^R = \left[\frac{\alpha_H Y_i^R}{(1 + t_{a(i)}) \bar{H}_i^R} \right]^{\frac{1}{1+\epsilon_{R,i}}}, \quad H_j^C = \sum_{g \in \mathcal{G}} H_{jg}^C. \tag{27}$$

Because the entrant's location-sector demands floor space in proportion to its output, entry competes with other producers for commercial space.

Residential land is owned by the CZ's residents. Rental income $\mathcal{R} = \sum_{i \in \mathcal{J}} r_i^R H_i^R$ is paid out through a local portfolio according to fixed housing-value shares ω_k^0, so each current resident of skill k receives

$$D_k = \frac{\omega_k^0 \mathcal{R}}{N_k}, \tag{28}$$

where N_k counts all current residents of skill k, including renters, nonworkers, and new arrivals. Commercial rents accrue to absentee landlords.

4.5 Government

Each county provides public services in proportion to the number of residents and developed land within its jurisdiction. Let N_a denote the number of modeled residents in county a and $\mathcal{J}_a$ its modeled tracts. I measure developed property using residential and commercial space valued at baseline rents,

$$\mathcal{H}_a = \sum_{j \in \mathcal{J}_a} \left(r_j^{R,0} H_j^R + r_j^{C,0} H_j^C \right). \tag{29}$$

County spending on public services is

$$G_a = G_a^0 \left[\varsigma \frac{N_a}{N_a^0} + (1 - \varsigma) \frac{\mathcal{H}_a}{\mathcal{H}_a^0} \right], \tag{30}$$

where ς is the resident-service share of baseline spending. The county finances this spending with the property tax in (25), levied on annual pretax rent. Its tax rate t_a adjusts to balance its budget,

$$\mathcal{T}_a = t_a \left(\mathcal{R}_a + \mathcal{R}_a^C \right) = G_a, \tag{31}$$

where $\mathcal{R}_a = \sum_{i \in \mathcal{J}_a} r_i^R H_i^R$ and $\mathcal{R}_a^C = \sum_{j \in \mathcal{J}_a} r_j^C H_j^C$ are the county's residential and commercial rental income. Baseline spending equals baseline property-tax receipts.

The county produces public services with goods bought outside the CZ at a fixed price, so government spending does not raise demand for local service firms and is not paid out to households. Because spending grows with residents and developed land, services per resident and per unit of developed land stay at their baseline levels, and their value to residents is part of the baseline amenities. Appendix B.9 gives the full property and government accounts.

4.6 Equilibrium and the Effects of Entry

Given the parameters in Table C.1, the national populations $\mathcal{N}_k$ and utility elsewhere in the country, and the location fundamentals (productivities A_{jg}, amenities and preference parameters, floor-space shifters $\bar{H}_i^R$ and $\bar{H}_j^C$, fixed-cost scales φ_{jg}, travel times t_{ij}, and baseline county spending G_a^0), an equilibrium is a set of wages w_{jk}^g, rents r_i^R and r_j^C, county tax rates t_a, service prices p_{jg}, establishment masses m_{jg}, and allocations of individuals N_k, N_{gk}, E_{gk}, F_{ijgk}, and U_{ik} such that

1. individuals choose residence, sector, participation, and commute as in (9), (8), (10), and (11);
2. tradable producers break even, (23), and employ workers according to (24);
3. service firms price at a markup and enter until profits are zero, (18)–(20), and employ workers according to (21);

4. labor markets clear at each workplace, sector, and skill, so the commuting flows from all feasible residences, including inbound commuters, equal employment,

$$\sum_i F_{ijgk} = L^g_{jk}, \qquad g \in \mathcal{G}, \quad j \in \mathcal{J}; \tag{32}$$

5. residential and commercial floor-space markets clear, (27);

6. service revenue in each tract and sector equals the spending directed to it, $R_{jg} = \sum_i \mathcal{E}_i \pi^I_{ijg}$;

7. residential rental income is paid out through the portfolio, (28); and

8. each county's tax rate balances its budget, (31).

The tradable price and the price of intermediates are set nationally and utility elsewhere in the country is fixed, so the CZ is a small open economy. Appendix B.10 writes out the full system and the feasible choices.

Each workplace tract has its own labor market for each sector and skill, and these markets are linked through individuals' choices, as in quantitative urban models with workplace-specific wages (Ahlfeldt et al., 2015; Monte et al., 2018). A higher wage at one workplace draws commuters from other workplaces, workers from other sectors, nonworkers, and migrants, so the employers they leave must respond.

Entry raises the productivity $A_{j_E g_E}$ of the entrant's tract and sector by the factor that raises employment there by J^E. I compare the equilibrium with this higher productivity to the equilibrium without entry, holding fixed all other productivities, amenities, preference parameters, housing expenditure shares, ownership shares, national populations, and the taste distribution. Wages, rents, floor space, prices, individual choices, service entry, and dividends differ between the two equilibria. I solve both in changes relative to the observed pre-entry economy (Section 5).

5 Quantification

I quantify the model separately for each event in its commuting zone (CZ), indexed by c, and its event year. The main parameters are the worker choice elasticities, migration, sector, inbound commuting, and participation, which jointly determine how readily the CZ attracts workers not previously tied to its labor market. The model implies gravity equations for commuting and shopping flows between tract pairs, and I estimate the commuting and shopping elasticities from their travel-time coefficients by PPML with origin and destination fixed effects. I estimate the sector and migration elasticities from shift-share regressions on public employment, earnings, and population data that vary across commuting zones. I calibrate the participation and inbound commuting elasticities to match the hiring sources in Fact 1. The remaining cost and supply parameters are set from BEA national accounts, shopping trip data, and the literature. Given these parameters (Table C.1), I recover the remaining primitives (amenities, productivities, fixed

costs, and the service spending share) by inversion so that the model reproduces each commuting zone's economy before entry. None of these steps uses the service responses of Fact 3, which leaves those estimates free to validate the model in Section 5.4.

5.1 Individuals

Sector choice. Individuals' sector choice determines whether the increase in tradable labor demand is met by workers in the same sector or workers who leave other sectors, governed by the elasticity ν. I estimate ν from a shift-share regression using data from the Quarterly Workforce Indicators. Across commuting-zone–sector–skill groups, I regress the change in log private employment on the change in log average earnings. Because earnings and employment are jointly determined, I instrument earnings growth with within-sector industry employment shares interacted with subsequent industry employment growth outside the CZ. The weighted two-stage least squares estimate is 1.64, with a CZ-clustered standard error of 0.51 and a first-stage F of 15.03, so a 1 percent rise in one sector's earnings relative to another's raises relative employment by about 1.64 percent over seven years, the horizon of my event studies. Appendix C.2.1 provides the full specification, the instrument, and the robustness checks. This regression is the empirical counterpart of the sector choice in equation (8). Among low-skill residents of a CZ, the number choosing one sector relative to another has elasticity ν with respect to the sectors' relative value, and the CZ-by-skill effects absorb the terms common to sectors. The mapping is approximate: average earnings stand in for the sector's value, and employment stands in for the number of individuals attached to the sector, so the coefficient also reflects the participation response, and I treat it as a benchmark for ν (Appendix C.2.1).

Migration. The migration elasticity μ_k determines how many workers the CZ draws from the rest of the country when the value of living there rises. I estimate it separately by skill with the same kind of shift-share variation used for sector choice. For each skill, I regress the 2010–2018 change in the log odds of living in a commuting zone rather than elsewhere in the country on the change in the value of living there, measured as log expected earnings (the employment rate times average earnings) less one quarter of log rent, with Census-division effects. I instrument the value change with the CZ's 2010 industry mix of skill-specific employment interacted with national industry employment growth outside the CZ. The estimates are $\mu_H = 2.32$ for high-skill and $\mu_L = 1.51$ for low-skill workers, with CZ-clustered standard errors of 0.75 and 0.51. High-skill workers are therefore more willing to move for better employment opportunities, consistent with the Current Population Survey evidence presented in Appendix A.3 that college graduates are about twice as likely as other workers to move for a new job (Appendix Table A.10).

Participation and inbound commuting. I calibrate the remaining two worker choice elasticities to the hiring sources in Fact 1, holding sector choice and migration at their estimates. The calibration chooses the participation elasticity η and three border ratios (q_S, q_M, q_L), which scale

each event's commuting elasticity into its border elasticity, $\psi_e = q_{\text{size}(e)}\chi_{c(e)}$. Participation is disciplined by the share of continuing-resident hires with no observed prior employment, which is 9.5 percent in the LEHD data. Some of these hires were previously self-employed, and their earnings are missing from the data, so I allow up to five percentage points of missing employment and treat 4.5 to 9.5 percent as the admissible range. The calibrated model predicts 6.8 percent are unemployed, nicely inside that range. I choose the four parameters by minimum distance. For each candidate vector, I solve the entry equilibrium for a quarter of firm entry events, with each productivity shock set so that employment in the entrant's tract and sector rises by the promised jobs, and compute where the entrant's recruits come from. I pool the events with weights that reproduce the LEHD data hire counts in six cells defined by entrant skill type and CZ size, and form four moments: the share of continuing-resident hires with no prior employment, and the share of hires who are new inbound commuters in small, medium, and large CZs. The parameters minimize the sum of squared deviations of these moments from the data. Across all moments, the root-mean-squared error is 1.3 percentage points. The calibration gives $\eta = 3.32$ and $(q_S, q_M, q_L) = (0.67, 0.81, 1.00)$.

Residence-workplace pair. Within the commuting zone, residential and workplace choices determine how far a labor demand shock spreads and where additional income accrues. I estimate commuting gravity separately by CZ and year using annual LODES flows from 2007–2022 and a fixed matrix of driving times between tract centroids using driving routes from OpenStreetMap. Fact 2 in Section 3 presents the 2019 comparison with shopping flows. The commuting coefficient satisfies $\beta^C_{ct} = -\chi_{ct}\kappa$, so gravity identifies only the product of the choice scale and the per-minute cost. Following Ahlfeldt et al. (2015) and Tsivanidis (2026), I normalize the per-minute commuting cost to $\kappa = 0.01$, which gives $\chi_{ct} = -\widehat{\beta}^C_{ct}/0.01$. In the 526 CZs included in the 2019 comparison, χ has a median of 6.29 and ranges from 0.87 to 10.80.

Shopping. Where workers shop determines where the entrant's payroll becomes service revenue. With CES demand across service destinations and the same per-minute travel-cost parameter κ as commuting, the coefficient on travel time in a gravity equation for shopping visits is $\beta^S_c = -(\sigma_{N,c}-1)\kappa$, so the rate at which visits decay with travel time reveals how substitutable destinations are. I estimate β^S_c for each commuting zone from 2019 Veraset data restricted to the modeled service sectors (Appendix C.3.2 gives the estimating equation). With $\kappa = 0.01$ per minute, the estimates imply $\sigma_{N,c}$ between 3.95 and 18.37, with a median of 13.28. The same elasticity sets service markups and the share of revenue available to cover fixed costs under free entry.

Individual budgets. I calibrate the housing expenditure share $\alpha_{H,c}$ using household and person records from the American Community Survey (ACS) Public Use Microdata Sample (PUMS).[24] Housing expenditure is annual rent for renters and predicted rent for homeowners, estimated from renter households using dwelling characteristics and location effects. I set $\alpha_{H,c}$ to the weighted

[24] I use one-year ACS data where available for the housing shares, but I fall back on the 2008–2012 five-year ACS file if the one-year housing shares are not available for that geography.

median of housing expenditure divided by disposable income within each CZ, where weights combine ACS household survey weights with geographic allocation to the CZ. In the national comparison, the median housing expenditure share ranges from 13 to 35 percent, with a cross-CZ median of 20.2 percent. Figure A.11 compares the housing budget share of a typical household in each CZ with the share of the CZ's total household income devoted to housing. Appendix A.7 provides the measurement details.

After housing, individuals in CZ c spend the fraction $s_{N,c}$ of their remaining resources on services, common to both skills. I set $s_{N,c}$ in the inversion of Section 5.3, so that household spending covers the service revenue implied by each CZ's payroll. Between the 10th and 90th percentile event baselines, the calibrated shares range from 53 to 76 percent. Thus, at fixed prices, an additional dollar of disposable income generates $s_{N,c}(1-\alpha_{H,c})$ dollars of service expenditure. Workers split this spending between shopping trips that start from home and trips that start from work. I measure this split in 2024 Veraset data for the entire United States, assigning each service visit to an individual's most recent home or work stop, so that a visit in a trip chain is attributed to the stop where the chain began. Among visits whose origin I can classify, 74 percent start from home and 26 percent from work, so I set the home share of service spending to $\rho = 0.74$.

I calibrate nonemployment income y_k^U using ACS PUMS person records for individuals aged 25-64, taking average public assistance, Supplemental Security Income, and Social Security income among nonworkers of skill k, including those receiving no payments.

I construct the housing ownership shares ω_{ck}^0 from ACS PUMS housing records linked to the individual's education. The model allocates residential rental income across skills in proportion to these ownership-value shares and equally among individuals within each skill group. I let ownership differ by skill (Fajgelbaum and Gaubert, 2020) because high-skill residents hold a disproportionate share of housing value in the data.

5.2 Firms and Space

Production cost shares. Tradable and service firms have the same Cobb–Douglas production function in labor, commercial space, and intermediate inputs, whose exponents equal each input's share of costs. Intermediate inputs are supplied from outside the CZ at a fixed price. To set these cost shares, I measure purchased intermediate inputs, net of purchased real-estate services, in the 2012 BEA supply–use accounts and split the remaining share of gross output 80/20 between labor and commercial space, following the labor–floor-space calibration in Ahlfeldt et al. (2015), Hoelzlein (2023), and Miyauchi et al. (2025) (Appendix C.5.2). Intermediate inputs account for 0.34 of gross output in services, 0.39 in high-skill tradables, and 0.67 in manufacturing, which gives labor shares of 0.53, 0.49, and 0.27 and commercial-space shares of 0.13, 0.12, and 0.07. Because intermediate inputs are bought at a fixed price, this production function is equivalent to a value-added production function in labor and commercial space with shares of 0.8 and 0.2 in every sector. For tradables, including intermediate inputs changes only the entrant's productivity shock, which is a gross-output gain, a fraction $1-\alpha_X^g$ of the equivalent value-added gain. For services, revenue

is measured as sales, including the inputs that service firms buy, so the service spending share calibrated in Section 5.3 is comparable with the share of services in consumer expenditure data.

Skill weights. Each sector combines high- and low-skill workers with Cobb–Douglas weights $\lambda_{jgk} = \lambda_{gk}$, which are common across tracts and fixed in counterfactuals. I set them to each skill's national share of the sector's wage bill in the 2012 Quarterly Workforce Indicators. Within each sector, I compute each education group's annual wage bill as employment times average earnings and sum across states and industries. Workers with a bachelor's degree or more receive 55 percent of the wage bill in high-skill tradables, 26 percent in low-skill tradables, 50 percent in high-skill services, and 27 percent in low-skill services, and low-skill workers receive the rest. The weights fix each skill's share of the wage bill rather than its share of employment, so employment shares still respond to relative wages.[25]

Establishment entry. Free entry sets the number of service establishments in a tract with elasticity ϵ_m with respect to net revenue, and the fixed-cost shifter φ_j is recovered from the observed establishment count (equation (C.39)). I set $\epsilon_m = 0.40$, in line with a county regression of five-year growth in service establishments on growth in service revenue relative to wages. Establishments and payroll information come from the County Business Patterns, and wages come from the Quarterly Workforce Indicators. County Business Patterns does not report revenue, so I impute it as payroll divided by each sector's national labor share of value added in the BEA accounts. Because the shares are fixed over time, the division affects only the weight of each sector in county revenue. The OLS coefficient is 0.39 (standard error 0.01), so a 10 percent rise in net revenue adds about 4 percent to the number of establishments (Appendix C.5.3). Because entry also raises revenue by attracting spending, and the revenue measure moves closely with employment, I treat the value as a descriptive calibration rather than a causal estimate. The value is also close to the response implied by Almagro and Domínguez-Iino (2025) for Amsterdam, where congestion in commercial real estate makes the number of establishments rise with amenity spending with an elasticity of about 0.40. They calibrate the congestion parameter from the range of housing supply elasticities in Saiz (2010).

Housing and commercial space. Residential supply uses tract-specific elasticities $\epsilon_{R,i}$ mapped from Baum-Snow and Han (2024). I retain each valid nonnegative prediction and replace a missing or negative value with the unweighted mean of the eight nearest tracts with valid predictions. Distance is measured between 2020 tract centroids. This extends the existing estimates without using employment or payroll outcomes. Across the 36,557 distinct tracts covered by the current event baselines, the unweighted mean is 0.53 and the median is 0.57. I set the commercial space

[25] QWI education categories cover workers aged 25 and older, so these shares proxy for the wage bills of the model's labor bundles. The wage bills include private and state and local government jobs. Industries keep the high- and low-skill-intensity classification used throughout. Industries with missing or ambiguous labels are allocated using the national employment-weighted mix of classified industries within the same broad industry.

elasticity as $\epsilon_C = 0.60$, at the upper end of the metropolitan office estimates of Marcato and Tong (2025, p. 412), which range from 0 to 0.61, and above the retail estimate of 0.49 in Benjamin et al. (1998, p. 9). This is close to the median residential elasticity across the current tracts. Greaney et al. (2024, pp. 29, 52) take a similar approach. Noting the absence of local estimates for commercial floor space, they set commercial supply elasticities in their model of the San Francisco Bay Area from residential ones.

Property taxes. Users of housing and commercial space pay the property tax $t_c = 0.28$ on top of the rent. Because the model has annual rents but no property values, t_c is the tax bill per dollar of annual rent rather than per dollar of value. A tax of 1 percent of value on a property worth 20 years of rent is 20 percent of a year's rent. Couture et al. (2024, pp. 878–879) measure taxes the same way and find that they are 20 percent of the annual user cost of housing downtown and 30 percent in the suburbs. The 2010–2014 American Community Survey implies about 32 percent of annual rent.[26] I use an average of these estimates and apply the same rate to commercial space.

Government service costs. I set the resident-service share ς to 0.79, the share of resident services in county-government spending on resident and developed-property services in the 2012 Census of Governments. Resident services are health, welfare, police, corrections, libraries, and recreation. Developed-property services are roads, drainage, sewerage, waste collection, fire protection, and transportation infrastructure. The share is the same across counties, and each county's baseline spending G_a^0 equals its baseline property-tax receipts.

5.3 Constructing the Pre-Entry Economy

I recover the remaining fundamentals so that each event's baseline equilibrium reproduces its pre-entry economy. Given observed wages, rents, and commuting costs, residential amenities B_{ik}^R and workplace–sector amenities B_{jgk}^W match residence and workplace employment by skill. Tracts that have more residents than their rents and access to jobs predict must offer high amenities. The same logic applies for workplaces that employ more workers than their wages predict. The taste shifters $\overline{B}_k^R$, $\overline{B}_{gk}^S$, $\overline{B}_{gk}^E$, and $\overline{B}_{gk}^U$ match the CZ's population, sector shares, and participation rates by skill, and tradable productivity sets baseline unit costs equal to the national price.

Service productivities are recovered so that the model reproduces where service purchases occur, adapting the procedure of Oh and Seo (2023). I infer each tract's service revenue from its worker payroll, which equals the labor share of costs under free entry. I then set the service share of individuals' nonhousing spending, $s_{N,c}$, so that the spending of the CZ's residents covers total CZ service revenue. Given this total budget, I invert service productivities so that consumers' demand, which follows the gravity structure of shopping, matches the service revenue observed in each tract.

[26] Because t_c is levied on top of the rent, taxes are $t_c/(1 + t_c) \approx 22$ percent of what users pay, which lies within this range.

The fixed-cost parameter then makes revenue per establishment consistent with free entry at the observed establishment count.

The calibrated service spending shares have a median of 0.63 and a range from 0.53 to 0.76 between the 10th and 90th percentiles of CZs. They are above the 0.48 share of nonhousing consumer spending that goes to local services as measured in BEA data. The residents of the CZ in the model are all workers observed in the LODES data plus the number of working-age nonemployed residents from the ACS data. This does not cover all of the people living in the CZ nor does it capture the full demand for local services. In the 2008–2012 ACS, for every 100 residents of the CZs who are either wage and salary employees or nonemployed and of working age (25 to 64), there are 32 other residents who report income but are not captured by the model. The majority are retirees (nonemployed residents aged 65 and over). Counting them, the population is therefore 1.32 times the population in the model, which matches the ratio of the median calibrated spending share to the observed BEA share, $0.63/0.48 = 1.31$.[27] The spending share of the model's individuals must support the service demand of these other consumers. I assume that these consumers have the same preferences as the model's individuals, so they spend the same share of their budgets on local services and choose where to shop in the same way. Thus, the choices of the model's individuals approximate theirs, and the calibrated $s_{N,c}$ adds their demand as a fixed multiple of the demand of the model's individuals, before and after entry.

I recover the amenity, productivity, fixed-cost, and taste fundamentals to rationalize the observed baseline economy as an equilibrium. These fundamentals are held fixed across counterfactuals. I solve for the entry equilibrium in changes relative to the baseline using exact hat algebra. Writing $\widehat{x} = x^1/x^0$, the fundamentals cancel, so the equilibrium conditions depend only on baseline shares and quantities, the parameters described above, and the entry shock. I solve jointly for the changes in wages, rents, service establishment counts, and dividends that clear all markets.

5.4 Validation

The purpose of this model is to explain the spatial distribution of the service multiplier. The model is successful only if it correctly predicts where the entrant's workers live, how much of its payroll becomes new income for CZ residents, and how that income translates into nearby service jobs. I exploit the fact that estimation makes limited use of the entry events. Only the participation and border elasticities are calibrated using one-fourth of the entrants from the public data to match the hiring shares from the Census events. The shopping elasticity comes from 2019 trip data, after every event, but it fixes only how fast trips decay with travel time, not the model's shopping flows. The model inversion uses data in the year before entry to reproduce the baseline economy of each commuting zone.

[27] For every 100 residents in the two groups the model covers, 31.6 others report positive personal income. Of these, 17.7 are nonemployed and aged 65 or over, 7.1 are self-employed, 3.6 are nonemployed and aged 18 to 24, and the rest are mostly nonemployed residents aged 15 to 17 and military personnel. The population ratio is $(100 + 31.6)/100 = 1.32$, and the ratio of spending shares is $0.631/0.48 = 1.31$.

Who is hired. First, the out-of-sample fit on the recruitment sources in the model closely mirrors the shares observed in the Census events. On the 161 firm entry events not used in the calibration of the participation or in-bound CZ share of workers, the model's split of hires among continuing residents, migrants, and outside-CZ commuters is 57, 18, and 25 percent against 58, 14, and 28 in the Census. Additionally, the income of the productivity shock matched to the entrant's hiring goes to the same workers in similar proportions as the Census data. Table 4 presents the share of each $100 of entrant payroll that becomes new income for residents. On the held-out events the model's shock generates the same net injection of about a third of entrant payroll, with no entrant wage premium and no moment that fixes the gain of job switchers. The income gains being close is more of a statement about the hiring shares and wages matching. The model pays each hire the wage of the tract and sector it joins, and the hire-weighted average of that wage is $102,000, between the data's median of $75,300 and mean of $125,000 for the same hires.

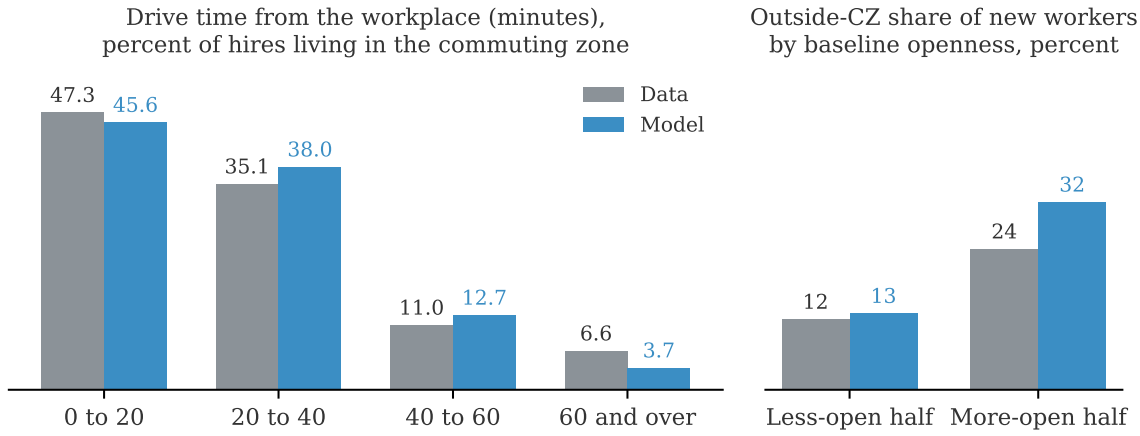
Table 4: The hired cohort on the events outside the calibration sample: payroll per $100 of entrant payroll (dollars) and residence within the commuting zone

	Census data	Model (not in calibration)
Paid to workers living outside the CZ	28	25.1
Replaces residents' prior earnings	40	39.7
New income of residents	32	35.2
migrants	15	18.8
previously nonemployed	4	3.3
job switchers' earnings gain	14	13.1
Hires living within 20 minutes of the site (percent of resident hires)	47.3	45.6
Events	400	161

Notes: The Census values are the Fact 1 earnings approximations for hires at large tradable entrants, rounded and without confidence intervals. The three components of residents' new income sum to it up to a $1 offset for out-movers, which is omitted. Model values are the workers who switch into the entrant's tract and industry, pooled over events with cohort-payroll weights. The residence row pools all 216 events with net-hire weights, commuting-zone residents only, since the model's outside area has no travel time; its Census value is the Fact 1 tabulation of hires' residences.

The model is able to closely match the distance of its workers' residential neighborhoods as shown in Figure 6. This indicates that the commuting gravity estimated on baseline flows predicts out of sample, without inverting bilateral amenities to match every tract pair's flow exactly (Dingel and Tintelnot, 2026). I then check whether the model pulls workers from outside the commuting zone in proportion to a site's prior openness to outside commuters. The border term is calibrated to the inward-commuting share of Census hires within three size groups, not to how that share varies with prior openness. I use the public event sample to divide the entrant sites at the median of prior openness and measure the share of new workers. This rises from 12 to 24 percent in the data and from 13 to 32 in the model. This is an important margin that separates a high-multiplier commuting zone from a low one.

Figure 6: Drive time from the workplace and the openness split, data and model (percent)



Notes: Grey bars are the data, blue bars the model. Left panel: hires at the entrant who live in the commuting zone by drive time from the workplace, as a share of all hires. The model reports the same share for the workers in the tract and sector of the shock. Right panel: the share of each subsidized site's new workers who live outside the commuting zone, public LODES at the 216 award sites split at the median of baseline commuting openness, against the model with promised-job weights. The drive-time bins are not targeted; the outside-CZ level is targeted within size groups and its split by openness is not.

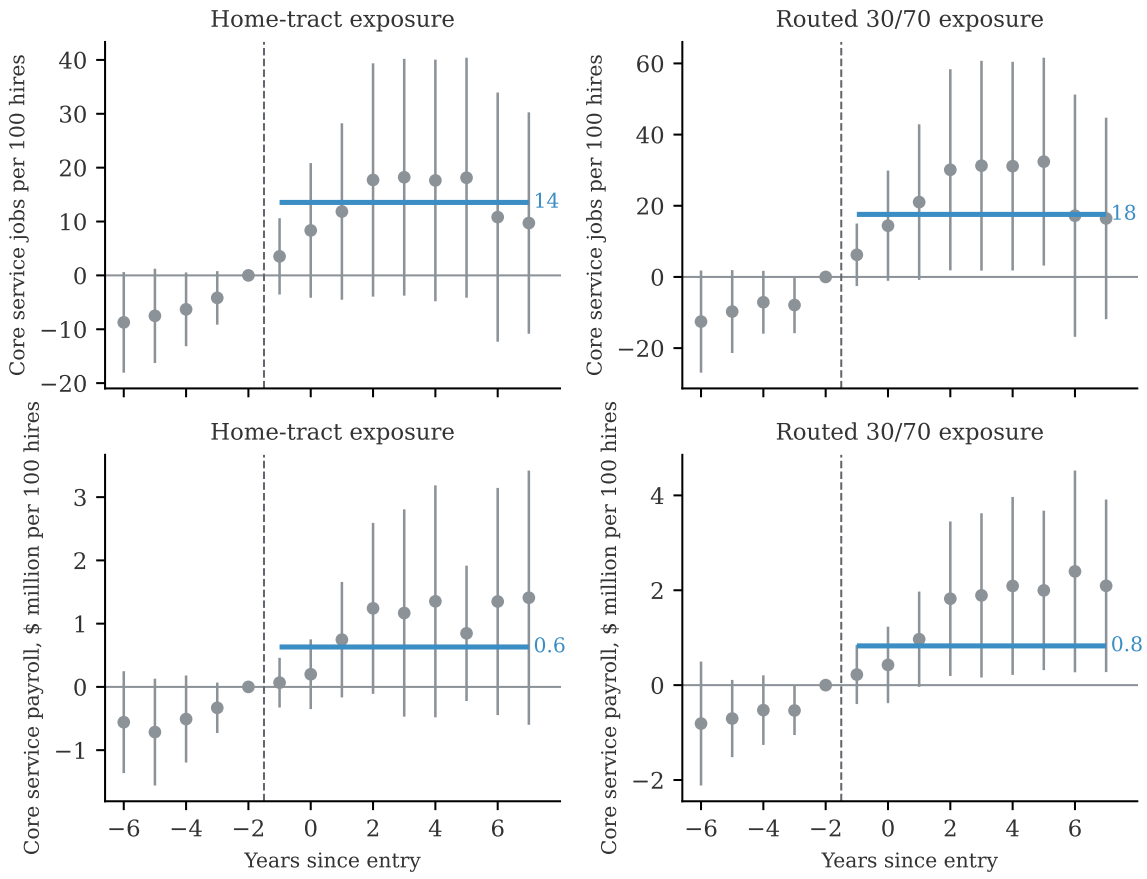
Where service jobs appear. Where service jobs appear in the model is broadly consistent with the patterns observed in the Census data. The central validation is the continuous neighborhood event study of Fact 3 for core services, retail, food and accommodation, and other services (NAICS 44 to 45, 72, and 81), which I replicate on the model for all firm entry events. I measure a tract's exposure as the number of the entrant's hires who live there and convert it into local demand for services either by claiming all of a person's spending occurs in their home tract, or giving 30 percent to their home tract and 70 percent spread over the adjacent tracts.[28] The outcome is the change in the tract's core service employment.[29] Figure 7 places the model beside the Census estimates. Under both exposures the model lies within about one clustered standard error for jobs and, for payroll, within one under the home-tract design and between one and one and a half under the routed design. The model does miss the spread more. Routing the hires 70 percent to the adjacent tracts raises the empirical coefficient by a factor of 1.7 and the model's by 1.3, in both jobs and payroll, so the model places too much spending in the home tract and cannot capture the concentration of services.

The model has two limits here. First, it performs much worse against the data when the service

[28]The 30/70 split follows the hyperlocal share of service visits in the 2019 Veraset data, the 11 percent of within-CZ core service visits that end in the home tract against the 24 percent that end in an adjacent tract (Table 3). The home tract and its neighbors carry the part of shopping that differs across tracts, while the rest of a resident's visits follow the common hub pattern of the commuting zone (Section 3), so the routed exposure captures the component of shopping that is specific to the shopping footprint of where the hire lives.

[29]Each event is solved against its own no-entry equilibrium with parameters and shocks fixed. The Census data estimates on approximately 100 commuting zones; the model's values are static event-fixed-effect projections on 216 events, so samples and estimators differ. The comparison for all nontradables, which adds education and health, is in Appendix D.

Figure 7: Core service employment and payroll where the entrant's workers live: empirical event study and the model



Notes: Circles are the continuous event-study coefficients of Fact 3 for core services (NAICS 44 to 45, 72, 81) per 100 exposed hires with 95 percent confidence intervals, with exposure in the hires' home tracts (left) and routed 30 percent to the home tract and 70 percent to the adjacent tracts (right); year −2 is the reference and hiring begins in year −1 (dashed line). The blue line is the model's response from the same design applied to its own hires on all 216 events. The model is static.

sector includes health and education. With all nontradables the model's response is about half of the empirical one in jobs and two fifths in payroll under the home-tract exposure, and two fifths and one fifth under the hyperlocal exposure (Appendix D). Schools and hospitals are funded in large part not by individuals' disposable income and they are located in specifically zoned places, neither of which the model is equipped to address. Second, the model cannot reproduce the stacked design that compares the top and bottom quartiles of exposed tracts. The gap in exposure between those quartiles is two to three times larger in the data than in the model, because hires' residences are concentrated in a few tracts near each site in the data and spread more evenly in the model. The two shortfalls compound to a contrast of a few jobs per 100 hires in the model against several dozen in the data.

6 Employment and Spatial Incidence

I use the quantified model to compare the baseline economy with the economy after a productivity increase in the entrant's tract and tradable sector calibrated to increase employment by the number of new jobs the entrant promises. I quantify the model for 216 firm entry events in 95 commuting zones. This sample of firm entry events comes from the subsidy awards data described in Section 2 and Appendix A.1.

The firm entry (productivity) shock adds about four fifths of the jobs it promises. For every 100 entrant jobs, the median event adds 80 jobs to the commuting zone. Increased demand for services translates into increased employment of 29 jobs in the service sector. Tradable employment outside the shocked location-sector falls by approximately 49 jobs. In the model, tradable sectors sell at fixed national prices. Tradable producers compete for workers and floor space but sell at a national price, so they cannot pass higher costs on to consumers and shrink instead. They shrink not only because of increased labor demand from the entrant, but also because of increased labor demand by service firms. Measured per net tradable job, the median service-job gain is 58 per 100.

Counting all jobs, the total multiplier is about 0.8 local jobs per entrant job, below the range of 1.3 to 1.8 that Bartik and Sotherland (2019) estimate for commuting zones. The service-job gain alone is about a third of the 160 in Moretti (2010). My model takes a narrow view of services and excludes construction, professional services, and business services that could all reasonably depend on local spending, in contrast to these benchmarks, which include everything outside manufacturing in Moretti (2010) and jobs in every industry in Bartik and Sotherland (2019). The model also leaves out two channels that would raise the total: Firms buy their intermediate inputs from outside the CZ, so there are no supply-side effects, and the entrant generates no productivity spillovers for other tradable producers, which would prevent some of the tradable job loss.

Following $100 of additional payroll in the shocked tract and industry, total CZ workplace payroll rises by about $81, after accounting for employment and wage changes elsewhere. Wage earnings rise by about $41 for residents of the CZ. The difference reflects earnings paid to workers living outside the CZ and changes in residents' earnings for residents who previously worked outside the CZ (9 percent of hires in the data, Appendix Table A.5). After accounting for dividends from ownership of residential land, housing costs and spending on tradable goods, spending at service firms inside the CZ rises by about $31. At the median, residents' welfare rises by about $11.9 per resident a year, or about $7.9 million across the CZ.

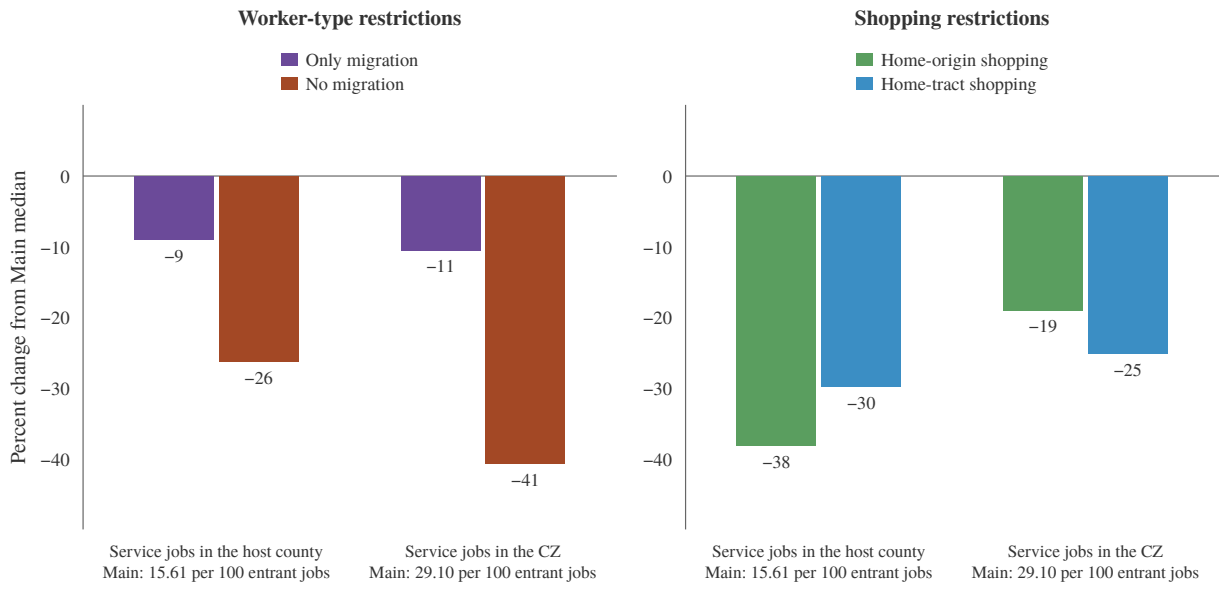
Section 6.1 shows what shopping, commuting, and migration add and traces the spatial incidence across neighborhoods. Section 6.2 asks how much the multiplier varies across markets and across locations within a county. Section 7 asks what each job costs the county that subsidizes it.

6.1 The role of commuting and shopping in the size of the multiplier

The model departs from other spatial models by allowing individuals to travel to consume services and by allowing them to migrate, live outside the CZ, and not participate in the workforce. Fact

1 showed that the in-bound commuting, migration, and participation margins are important for matching the shares of workers hired by these entering firms, and they distinguish new workers from new service consumers. Fact 2 showed that allowing travel for services is important for models with small units of location because only 10 percent of service trips occur in the tract where individuals live. In Figure 8, I keep the same productivity shock, and I shut down one feature of the model at a time. The left panel reports the two restrictions where I close a source of the entrant's workers either by fixing the regional population, or by shutting down in-bound commuting and fixing participation so all new workers are migrants. The right panel shows restrictions on service trips where I require all trips to originate at home or originate and end in individuals' home tracts.

Figure 8: Percentage changes in host-county and CZ service multipliers under four restrictions



Notes: Bars report the percent change of a restriction's median service multiplier from the main-model median. Every specification subtracts its own no-entry baseline. Home-origin shopping retains endogenous destinations; home-tract shopping uses the nearest supported tract when necessary. No migration fixes each individual's residence region while retaining participation and cross-border commuting. Only migration removes nonworkers and cross-border commuting in both equilibria.

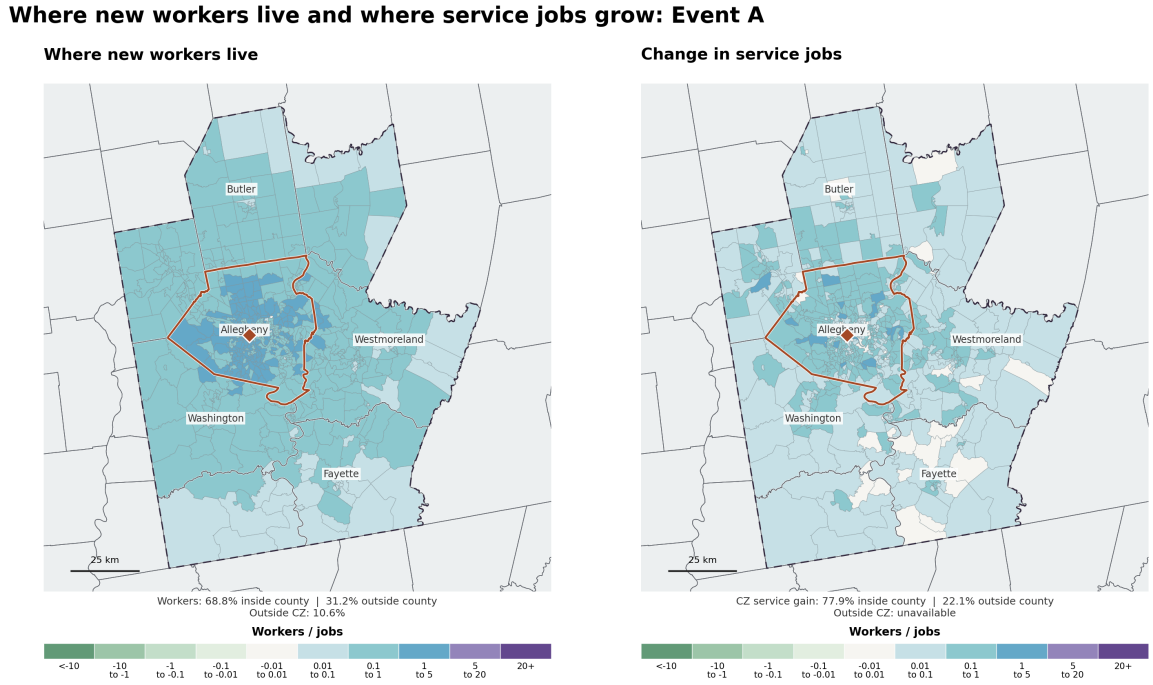
The left panel of Figure 8 shows that restricting either source of the entrant's workers lowers the service multiplier. When migration is the only source of new workers, the multiplier falls even though service revenue rises by more than in the main specification with commuting. Without participation or cross-border commuting, the entrant can draw on workers outside the CZ's existing workforce only if they move to the CZ. Moving is costlier than commuting or entering the labor force, since a migrant must leave their current residential location and compete for housing in the CZ, so the pool of new workers is smaller. The entrant fills more of its jobs by bidding workers away from other local employers. Tradable job losses at other employers rise from 49 to 81 per 100 promised jobs, and service wages rise about twice as much as in the main model. In this situation,

higher spending supports fewer additional service workers because more revenue is absorbed by higher wages for the existing workforce. Without migration, jobs are filled by existing residents and commuters, so less of the entrant's payroll becomes new resident income and, as a result, service spending. Migration on its own supplies more new service spending per worker, while participation and commuting ease competition for labor.

The right panel shows that the multiplier also falls when I restrict how individuals can travel for services. When every service trip must start from home, the multiplier falls about twice as much in the host county as in the CZ. Trips that start at a workplace are 26 percent of all shopping trips, but they account for a majority of the increase in service revenue in the host county while only accounting for 13 percent in the rest of the CZ. For each $100 of payroll the entrant adds, payroll at host-county workplaces rises by about $95, but the earnings of host-county residents rise by only $26, because more than half of the entrant's workers live in another county. Residents' incomes and service budgets are almost the same as in the main model, but the spending lands in different places. A model that does not account for workplace origin shopping understates the multiplier close to the entrant. When shopping is restricted to the home tract, the county and CZ multipliers fall by similar amounts. The county loses the workers' spending but keeps its own residents' spending, which can no longer leak to neighboring counties. The multiplier declines because service firms near the entrant compete for limited commercial space, and pass the higher rents into prices: commercial rents in the entrant's tract rise by 3.7 percent, against 2.5 percent in the main model. This congestion makes living near the entrant less attractive, so fewer workers move into the CZ, more of the jobs are filled by commuters from outside, and less of the payroll becomes resident income.

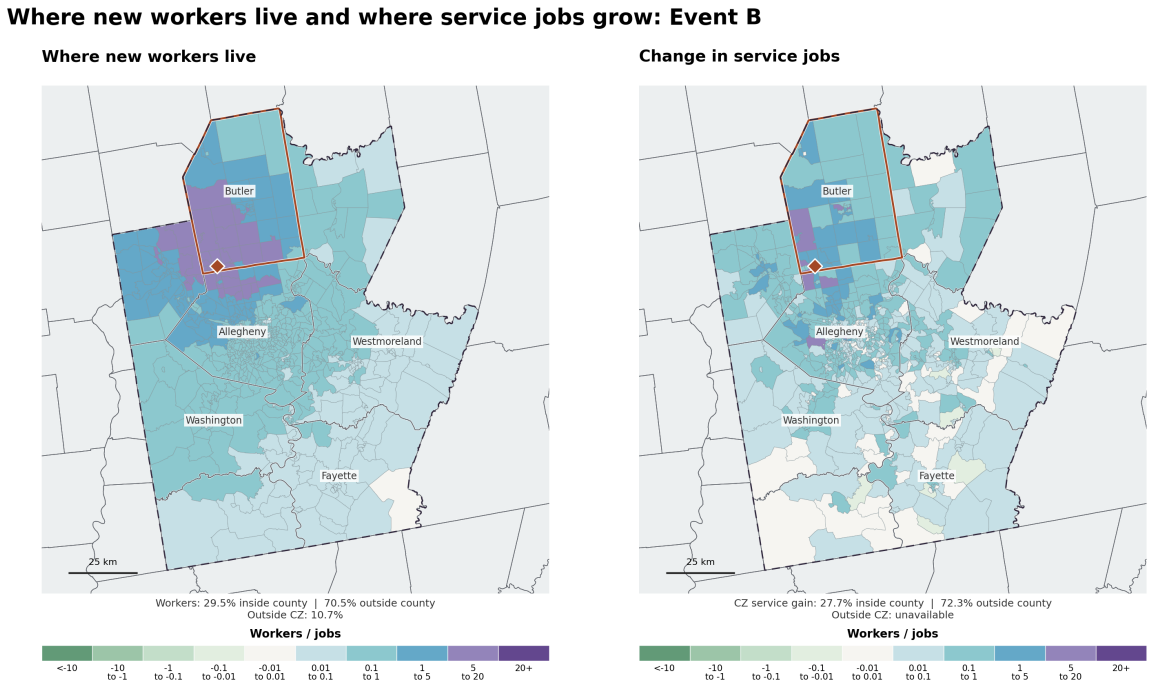
At the median, the county experiencing the productivity shock contains 28 percent of the CZ's residents but receives 61 percent of the resident income gain and 62 percent of the service job gain. Figures 9 and 10 show two entries in the same CZ that add about 31 and 32 service jobs per 100 entrant jobs in the CZ, respectively, but their counties see very different gains. Event A enters the CZ's central county, where most of its workers live, and that county keeps 78 percent of the gain. Event B enters at the edge of a neighboring county, where about 70 percent of its workers live outside the county, and it keeps only 28 percent, with most of the gain across the county line.

Figure 9: Where new workers live and where service jobs grow: Event A



Notes: Both figures show entry without a property-tax exemption. The left panel distributes the new jobs across residence tracts in proportion to the post-entry workforce in the entrant's tract and sector. The right panel shows the change in service jobs by workplace tract relative to no entry. Workers living outside the CZ are included in the reported shares but are not mapped. The diamond marks the entrant's tract and the orange outline marks its host county.

Figure 10: Where new workers live and where service jobs grow: Event B



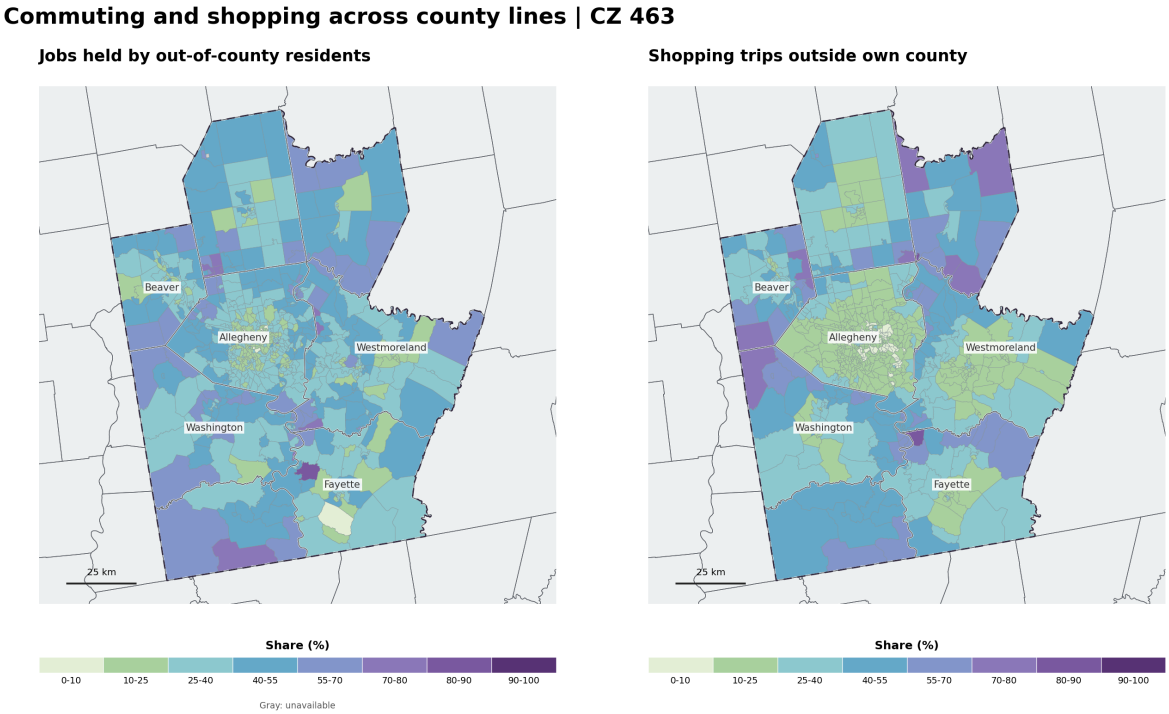
6.2 Variation in the multiplier

The medians reported mask the heterogeneity in the multiplier across productivity shocks. Between the 10th and 90th percentiles of firm entry shocks, the service gain ranges from about 18 to 39 jobs per 100 entrant jobs. The location of the shock within a county matters nearly as much as which CZ it enters. Within the typical host county, the tract at the 90th percentile adds about 1.7 times as many service jobs as the tract at the 10th, against a ratio of 2.2 across markets. Figure 11 shows how much tracts can differ even within one county. In Table 5, the service multiplier is lower when a firm enters a tract where more of the prior workers live outside the county and where more of their service spending leaves it. These maps show that these two leakage shares are positively correlated with distance to the border but overall shopping is more contained in the county.

A natural first explanation is that the productivity shock is calibrated to meet job targets that differ greatly in size, from 500 promised jobs at the 10th percentile to about 2,100 at the 90th. To control for this, I recalibrate the productivity shock at each of the 216 observed sites in each tradable industry, so that employment there rises by 500 jobs. I repeat this exercise so that annual payroll there rises by $70 million, the median amount promised. Fixing every shock at 500 jobs reduces the variance of the service gain by only a quarter.

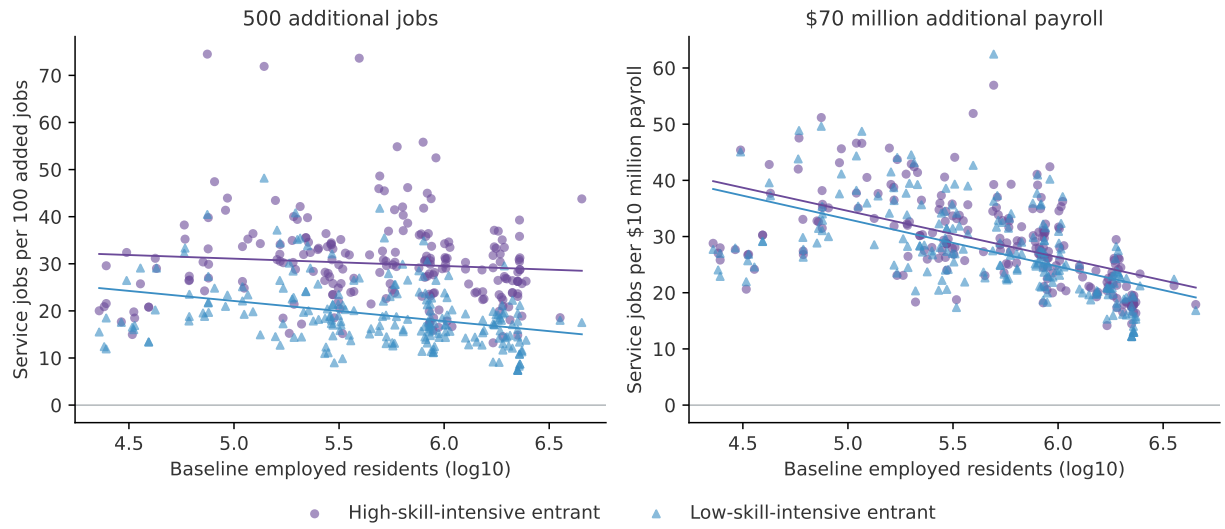
Figure 12 shows that market size barely moves the multiplier when every site receives 500 jobs, but lowers it steadily when every site receives $70 million of payroll. The line slopes down under the payroll shock because a dollar of payroll supports fewer service jobs in larger markets, for two reasons: Higher wages mean that each service job takes more revenue to support, about $130,000 in

Figure 11: Commuting and shopping across county lines in the CZ of Events A and B



Notes: CZ 463. The left panel shows the share of each tract's jobs held by workers who live in another county. The right panel shows the share of each tract's residents' shopping trips that end outside their own county. Denominators include all primary workplace jobs and all resident-origin shopping trips, including flows that cross the CZ boundary. The dashed line marks the CZ boundary; gray tracts are unavailable.

Figure 12: Market size and service responses under equal jobs and equal payroll



Notes: Each panel contains 432 shocks: both tradable industries at each of 216 observed sites. Lines are separate linear fits against log baseline employed residents. The left panel fixes the target tract–industry employment increase at 500 jobs; the right fixes its additional annual payroll at $70 million (2023 dollars), including wage changes on existing jobs. Productivity adjusts to each target. Blue triangles represent low-skill-intensive entrants and purple circles high-skill-intensive entrants.

the largest fifth of markets compared to $98,000 in the smallest, and this accounts for about three fifths of the decline per dollar of payroll. Housing is also less elastic, with a residential land supply elasticity of 0.46 in the largest fifth of markets against 0.70 in the smallest, so fewer workers move to large markets — 28 new residents per 100 added jobs against 39. Thus, less of the entrant's payroll becomes new resident income in larger markets, which accounts for the rest of the decline. Under the 500-job shock the lines are much flatter, because the same 500 jobs bring more payroll in a larger market, where wages are higher. Between the 10th and 90th percentiles of market size, the payroll added per job rises by about a third, while the service jobs supported by each dollar of payroll fall by about a third, for the two reasons above. More dollars per job and fewer service jobs per dollar nearly cancel, and the service gain falls only from 27 to 23 jobs per 100 added jobs. The offset is more evident for high-skill-intensive entrants, whose pay rises more with market size, than for low-skill-intensive entrants. The exercise shows that market size matters through local wages and housing supply, which a multiplier measured per job misses.

To further examine the determinants of the multiplier, I run regressions including county-level and location-level characteristics observed before entry. I regress the log of the host county's service jobs per 100 added jobs under the 500-job shock on these characteristics, with each coefficient scaled to a one standard deviation change in the variable. Tract characteristics alone explain 71 percent of the variation and county characteristics explain 44 percent (compared to 13 percent for CZ characteristics in the CZ multiplier (Table E.8, in the Appendix)).

Unsurprisingly, commuting and shopping leakage are the two largest determinants. The county keeps a service job only if new income translates to new spending inside the county. I measure commuting leakage as the share of the workers employed in the entrant's industry in the tract before entry who live outside the county. I measure shopping leakage as the share of those workers' service spending that leaves the county, routing each worker's service budget from home and from work with the baseline shopping shares. Both are higher near the county border, and their correlations with the tract's distance to the border are -0.58 and -0.45. In the quarter of tracts closest to the border, 67 percent of the workers live outside the county and 48 percent of their service spending leaves it, double and triple the amount leaving in the tracts farthest away. On their own, the two shares explain 62 percent of the variation in the multiplier.

However, the labor market in the tract matters as well. The multiplier is lower when there is a thicker labor market in the tract, or more of the same industry workers already employed there. In the model this means that tract is already productive for that industry, and the shock can recruit more workers with a lower payroll increase. Conversely, the multiplier is higher where the location pays more than the county average, because each added job then brings more payroll, about $131,000 in the quarter of tracts with the highest relative wage against $77,000 in the lowest, while the service jobs that this spending supports are paid at the county's wages, not the tract's. As shown in the figure above, county size does not matter in itself. Larger counties keep more service jobs in column (1), but this reflects less leakage from their sites, and county size loses significance once leakage is measured at the site. Among the county characteristics, both service

jobs per resident and the residential land supply elasticity raise the multiplier. A county with more service jobs per resident is where people already shop, and a county with more elastic residential land gains more new residents. In summary, the multiplier differs mainly because of the location, not the market. What matters is the site's place in two networks, where its workers live and where they shop, and not the size of the county. A model in which workers spend only where they live would capture only the first. Both networks are observed before entry, so a county can approximate its share of the gain before the firm arrives.

Table 5: Host-county and site characteristics and the host-county service multiplier

	(1)	(2)	(3)
Host-county characteristics			
Log baseline county residents	0.51***		0.06
	(0.09)		(0.06)
County service jobs per resident	0.24***		0.10***
	(0.05)		(0.03)
County tradable jobs per resident	−0.18***		0.05**
	(0.05)		(0.02)
Residential land supply elasticity	0.14***		0.09***
	(0.04)		(0.02)
Shopping elasticity, $\sigma_{N,c}$	−0.40***		−0.06
	(0.10)		(0.04)
Commuting elasticity, χ_{ce}	0.04		−0.02
	(0.08)		(0.03)
Service job concentration	0.11**		−0.01
	(0.06)		(0.03)
Site characteristics			
Log location-sector jobs		−0.11***	−0.12***
		(0.02)	(0.01)
Log location-sector wage relative to county average wage		0.21***	0.22***
		(0.01)	(0.01)
Share of site workers living outside the county		−0.09***	−0.15***
		(0.02)	(0.02)
Share of site workers' service spending leaving the county		−0.53***	−0.37***
		(0.03)	(0.03)
Industry effects	Yes	Yes	Yes
R^2 after removing effects (%)	43.8	70.6	74.6
Sites	2,214	2,214	2,214
Events	216	216	216
CZ clusters	95	95	95

Notes: The outcome is the log of host-county service jobs per 100 added jobs under the 500-job shock. The sample is each observed site and its alternatives in the host county, with the industry held at its observed type for 2,214 sites in 216 events, excluding 38 sites where the host-county service gain is not positive. County service and tradable jobs per resident are baseline county employment in each group over all residents. The residential land supply elasticity is the mean across the tracts of the sampled sites in the county. The shopping and commuting elasticities are estimated at the CZ level. Service job concentration is the sum of squared tract shares of the county's baseline service jobs. The relative wage is the log of the location-sector wage minus the log of the county's average wage across all jobs at baseline. The share of service spending leaving the county is computed in the no-entry equilibrium as the service budgets of the workers employed in the site's tract and industry routed from their homes and from the site using the baseline shopping shares. Parentheses give CZ-clustered standard errors; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, using two-sided t tests.

7 What Does a Job Cost, and Is the Subsidy Worth Paying?

The public data I collect include a promised job figure and a combined state and local subsidy amount for all firms. The simplest measure of what a job costs divides the total award by the number of jobs the firm promises. The median award is about $19,000 per promised job (2023 dollars). This is modest next to the discretionary deals over $5 million in Slattery and Zidar (2020,

Table 2), which average about $120,000 per promised job, and compares with about $9,900 per pledged job for the California Competes Tax Credit (Freedman et al., 2021). Promised jobs are not the jobs a place gains, and the adjustment can go either way. For California Competes, Freedman et al. (2021) estimate that each pledged job comes with about 2.5 jobs in the recipient tract, which lowers the cost to about $3,960 per job created.[30] In my events, total employment in the CZ rises by only about 80 jobs for every 100 promised, because other tradable employers lose workers, which raises the median award to about $23,000 per net job. Subsidies are paid out over multiple years, and the median contract length among the local subsidies I verify is ten years, so the awards in my data come to about $1,900 per promised job and $2,300 per net job per year.

None of these figures is the true cost of a job. These numbers overstate the fiscal cost because the entry brings new residents and workers who add income, sales, and property tax revenue. Additionally, they are not priced in welfare so a job count cannot speak to how much residents gain from the jobs or lose from the forgone tax. And they ignore the marginal decision, because the award buys the jobs only if the firm would not have entered without it. For a quarter of the awards in my data, I verify the actual amount provided by the county in which the firm locates. Because the model includes property taxes, it captures the higher property values that the entry brings, so the tax the county forgoes can be set against the revenue it gains from more residents and higher land rents. I represent the subsidy as an exemption on the property tax owed on the entrant's new commercial space and choose the exempt fraction so that the tax forgone in equilibrium matches the annual amount in the subsidy data.

At the median, the local subsidy forgoes about $2,200 a year per promised job and about $2,900 per net CZ job. If the service jobs are left out, an entry adds only about half of its promised jobs to the CZ, and the local subsidy costs about $4,700 a year per net CZ job. With the subsidy in place, the host county's revenue rises by about $5.5 million a year, of which about $2.5 million comes from the entrant and $3.0 million from the rest of the county. The new residents and floor space require about $1.9 million more in public services. The remainder, about $3.5 million a year or $4,200 per net CZ job, is returned to property owners through a lower tax rate, which falls in over 90 percent of scenarios, and by 11 basis points at the median.

Welfare per public dollar summarizes these gains and losses for residents, in the spirit of the marginal value of public funds (Hendren and Sprung-Keyser, 2020). When the subsidy brings the entry, residents across the CZ gain $3.1 for each dollar of forgone tax, against $1.1 for residents of the host county (Table 6). The host county bears the cost, while much of the gain goes to residents of the neighboring counties where the entrant's workers live and shop. These ratios heavily depend on the location of the firm entry within the county as described in the previous section.

A subsidy brings the gains of entry only when the firm would not have entered the locality without it. When the firm would have come anyway, the subsidy only loses tax revenue the county

[30] The Slattery and Zidar figures are in 2017 dollars; their text reports a mean of $160 million and $108,000 per job (pp. 91, 104). The Freedman et al. figures divide $1.23 billion in allocated credits by 124,000 pledged jobs and apply their tract-level multiplier of about 2.5 (Freedman et al., 2021, p. 24). The paper was later published as Freedman et al. (2023).

Table 6: Annual local-benefit costs and outcomes, conditional on entry

Measure	P10	Median	P90
Commuting zone			
Service jobs per 100 promised jobs	17.0	27.7	42.5
Annual tax expenditure per net CZ job ($)	178	2,870	8,372
Annual welfare per baseline resident ($)	5.7	17.0	60.6
Resident welfare per public dollar	1.1	3.1	44.4
Host county			
Service jobs per 100 promised jobs	4.0	17.2	32.7
Annual tax expenditure per net host-county job ($)	163	2,575	6,501
Annual welfare per baseline resident ($)	3.7	25.0	64.6
Resident welfare per public dollar	0.5	1.1	11.9

Notes: Equal-weight quantiles across 54 entry scenarios with a verified local subsidy. This excludes the state share of the package. The same local cost is divided by CZ or host-county net job gains. Welfare is annual residential equivalent variation relative to no entry. Ratios are calculated within each scenario before taking quantiles. All dollars are in 2023 prices.

would otherwise have collected. Bartik (2018) places the share of incentivized entries that typical incentives actually cause between 2 and 25 percent.[31] The local subsidy involves a trade-off between securing the firm entries that would otherwise not have happened and giving up tax on the entries that would have happened anyway. In the median scenario, CZ residents gain $3.1 for each public dollar when the subsidy causes the full entry (Table 6). When the firm would have come anyway, they lose about $0.43 for each dollar, because the county makes up the forgone revenue with a higher tax rate on other property. The subsidy breaks even for CZ residents when the gain on the entries it causes equals the loss on the rest, which happens when it causes about 12 percent of the entries it subsidizes. This threshold falls inside Bartik's range, so whether these subsidies pay off for residents depends on how often they actually change where firms locate (Table 7).

Table 7: Residents' expected gain from offering the subsidy, by the share of entries it causes

Share of entries the subsidy causes	100%	25%	2%
CZ residents' expected gain per subsidy dollar	3.06	0.46	−0.32
Scenarios where CZ residents gain	54	44	8
Host-county residents' expected gain per subsidy dollar	1.09	0.12	−0.17
Scenarios where host-county residents gain	52	38	7
Expected net CZ jobs per 100 promised jobs	76.4	17.5	−0.8

Notes: Medians and counts across the 54 recorded-subsidy scenarios in Table 6. Gains are per dollar of the subsidy's annual tax expenditure, and Appendix F.6 gives the calculation. County budgets balance in every scenario, so welfare already includes the taxes residents pay and the services they receive. The 25 and 2 percent columns bracket the range Bartik (2018) considers plausible. Welfare is annual residential-opportunity equivalent variation. Annual 2023 dollars.

[31] The lower end reflects the median of the estimates without an obvious bias, 3.4 percent, and the upper end is close to the 23.5 percent median of estimates biased upward (Bartik, 2018, pp. 1–2).

8 Conclusion

State and local governments spend over $20 billion a year on firm-specific subsidies, justified by the direct jobs these firms create and the indirect jobs their payroll supports. This paper asks how commuting and shopping shape the size of the local service multiplier and where it lands, and how the multiplier changes the cost-effectiveness of these subsidies. It follows the payroll from the firm to its workers, from the workers to the neighborhoods where they live and shop, and from those neighborhoods to local service firms, using linked worker histories and a tract-level model of an open commuting zone.

The data show that a new job is not the same as a new local consumer. Only half of an entrant's hires previously worked in the commuting zone, and the rest live outside it, move in, or were unemployed or self-employed. Each group brings a different amount of new income, and at least $32 of every $100 of entrant payroll becomes new income for residents. Residents spend that income close to home, and service employment grows in the neighborhoods where the entrant's workers lived before entry.

The model turns these pieces into a net effect. For every 100 jobs an entrant promises, the median commuting zone gains about 29 service jobs, or about 58 per 100 net tradable jobs, about a third of the 160 in Moretti (2010). It gains about 80 jobs in total once the jobs lost at other tradable firms are counted. Residents' welfare rises by about $11.9 a year for every resident of the CZ, or about $7.9 million in total. Commuting and shopping shape both the size of these gains and where they land. Migrants bring the most new spending per worker, while commuters and residents who enter the labor force bring less but ease the competition for workers. More than half of an entrant's workers live outside its county and shop mostly near home, so shopping trips that start at the workplace account for most of the host county's gain in service revenue.

The multiplier differs widely across entries, from about 18 to 39 service jobs per 100 entrant jobs between the 10th and 90th percentile events. The size of the entry and of the market explain little of this. Fixing every entry at 500 jobs removes only about a quarter of the variance, and market characteristics explain only about 13 percent, because a larger market's higher pay per job is offset by higher service wages and less elastic housing. Where in a county the firm locates matters nearly as much as which market it enters: within the typical county, the site at the 90th percentile adds about 1.7 times as many service jobs as the site at the 10th, against 2.2 across markets. The main reason is where the site's workers live and shop. Sites near the county border lose more than twice as many service jobs to neighboring counties as sites far from it. Both networks are observed before entry, so a county can approximate its share of the gain before the firm arrives.

The service multiplier lowers the cost of a job, but the county that pays keeps only part of the gain. For the 54 entries with a verified local subsidy, the subsidy forgoes a median of about $2,900 a year per net CZ job, against about $4,700 if the service jobs are left out. The host county holds 28 percent of the CZ's residents and receives 62 percent of its service-job gain. Median resident welfare per public dollar is $3.1 for the CZ and $1.1 for the host county, which bears the entire cost.

The model takes the firm's entry as given, so it cannot say whether a subsidy changed the location decision. When the firm would have entered regardless, the county simply forfeits tax revenue. If a subsidy causes 25 percent of the entries it subsidizes, CZ residents can expect to gain about $0.46 per subsidy dollar; at 2 percent, they can expect to lose about $0.32. The subsidy pays off for CZ residents only if it causes at least 12 percent of the entries it subsidizes, and for host-county residents at least 14 percent, both inside the range in Bartik (2018). Within that limit, the lesson is that what an entry is worth to a locality depends less on how many jobs the firm promises than on who fills them, where they live, and where they shop.

References

Abowd, John M., and Lars Vilhuber. 2011. "National Estimates of Gross Employment and Job Flows from the Quarterly Workforce Indicators with Demographic and Industry Detail." *Journal of Econometrics* 161 (1): 82–99.

Acemoglu, Daron, and Pascual Restrepo. 2020. "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy* 128 (6): 2188–2244.

Adão, Rodrigo, Michal Kolesár, and Eduardo Morales. 2019. "Shift-Share Designs: Theory and Inference." *Quarterly Journal of Economics* 134 (4): 1949–2010.

Agarwal, Sumit, J. Bradford Jensen, and Ferdinando Monte. 2020. "Consumer Mobility and the Local Structure of Consumption Industries." Working Paper 23616, National Bureau of Economic Research.

Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf. 2015. "The Economics of Density: Evidence from the Berlin Wall." *Econometrica* 83 (6): 2127–2189.

Almagro, Milena, and Tomás Domínguez-Iino. 2025. "Location Sorting and Endogenous Amenities: Evidence from Amsterdam." *Econometrica* 93 (3): 1031–1071.

Artuç, Erhan, Shubham Chaudhuri, and John McLaren. 2010. "Trade Shocks and Labor Adjustment: A Structural Empirical Approach." *American Economic Review* 100 (3): 1008–1045.

Autor, David, David Dorn, and Gordon H. Hanson. 2021. "On the Persistence of the China Shock." Working Paper 29401, National Bureau of Economic Research.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103 (6): 2121–2168.

Bartik, Timothy J. 1991. *Who Benefits from State and Local Economic Development Policies?* Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.

Bartik, Timothy J. 2018. "'But For' Percentages for Economic Development Incentives: What Percentage Estimates Are Plausible Based on the Research Literature?" Upjohn Institute Working Paper 18-289, W.E. Upjohn Institute for Employment Research.

Bartik, Timothy J., and Nathan Sotherland. 2019. "Local Job Multipliers in the United States: Variation with Local Characteristics and with High-Tech Shocks." Upjohn Institute Working Paper 19-301, W.E. Upjohn Institute for Employment Research.

Baum-Snow, Nathaniel, and Lu Han. 2024. "The Microgeography of Housing Supply." *Journal of Political Economy* 132 (6): 1897–1946.

Benjamin, John D., G. Donald Jud, and Daniel T. Winkler. 1998. "A Simultaneous Model and Empirical Test of the Demand and Supply of Retail Space." *Journal of Real Estate Research* 16 (1): 1–14.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2022. "Quasi-Experimental Shift-Share Research Designs." *Review of Economic Studies* 89 (1): 181–213.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2025. "A Practical Guide to Shift-Share Instruments." *Journal of Economic Perspectives* 39 (1): 181–204.

Busso, Matias, Jesse Gregory, and Patrick Kline. 2013. "Assessing the Incidence and Efficiency of a Prominent Place Based Policy." *American Economic Review* 103 (2): 897–947.

Caliendo, Lorenzo, Maximiliano Dvorkin, and Fernando Parro. 2019. "Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock." *Econometrica* 87 (3): 741–835.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna. forthcoming. "Difference-in-Differences with a Continuous Treatment." *American Economic Review*.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. "The Effect of Minimum Wages on Low-Wage Jobs." *Quarterly Journal of Economics* 134 (3): 1405–1454.

Couture, Victor, Cecile Gaubert, Jessie Handbury, and Erik Hurst. 2024. "Income Growth and the Distributional Effects of Urban Spatial Sorting." *Review of Economic Studies* 91 (2): 858–898.

Davis, Donald R., Jonathan I. Dingel, Joan Monras, and Eduardo Morales. 2019. "How Segregated Is Urban Consumption?" *Journal of Political Economy* 127 (4): 1684–1738.

Deshpande, Manasi, and Yue Li. 2019. "Who Is Screened Out? Application Costs and the Targeting of Disability Programs." *American Economic Journal: Economic Policy* 11 (4): 213–248.

van Dijk, Jasper Jacob. 2017. "Local Employment Multipliers in U.S. Cities." *Journal of Economic Geography* 17: 465–487.

Dingel, Jonathan I., and Felix Tintelnot. 2026. "Spatial Economics for Granular Settings." *Econometrica* 94 (2): 407–464.

Fajgelbaum, Pablo D., and Cecile Gaubert. 2020. "Optimal Spatial Policies, Geography, and Sorting." *Quarterly Journal of Economics* 135 (2): 959–1036.

Freedman, Matthew. 2013. "Targeted Business Incentives and Local Labor Markets." *Journal of Human Resources* 48 (2): 311–344.

Freedman, Matthew, Shantanu Khanna, and David Neumark. 2023. "Combining Rules and Discretion in Economic Development Policy: Evidence on the Impacts of the California Competes Tax Credit." *Journal of Public Economics* 217: 104777.

Freedman, Matthew, David Neumark, and Shantanu Khanna. 2021. "Combining Rules and Discretion in Economic Development Policy: Evidence on the Impacts of the California Competes Tax Credit." Working Paper 28594, National Bureau of Economic Research.

Fuest, Clemens, Andreas Peichl, and Sebastian Siegloch. 2018. "Do Higher Corporate Taxes Reduce Wages? Micro Evidence from Germany." *American Economic Review* 108 (2): 393–418.

Galle, Simon, Andrés Rodríguez-Clare, and Moises Yi. 2023. "Slicing the Pie: Quantifying the Aggregate and Distributional Effects of Trade." *Review of Economic Studies* 90 (1): 331–375.

Gardner, John. 2021. "Two-Stage Differences in Differences." Working Paper, arXiv:2207.05943.

Garin, Andrew. 2025. "Do Place-Based Industrial Interventions Help 'Left-Behind' Workers? Lessons from WWII and Beyond." Working Paper 33418, National Bureau of Economic Research.

Giuliano, Genevieve, and Kenneth A. Small. 1991. "Subcenters in the Los Angeles Region." *Regional Science and Urban Economics* 21 (2): 163–182.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2020. "Bartik Instruments: What, When, Why, and How." *American Economic Review* 110 (8): 2586–2624.

Greaney, Brian, Andrii Parkhomenko, and Stijn Van Nieuwerburgh. 2024. "Dynamic Urban Economics." Working Paper.

Greenstone, Michael, Richard Hornbeck, and Enrico Moretti. 2010. "Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings." *Journal of Political Economy* 118 (3): 536–598.

Hanson, Andrew, and Shawn Rohlin. 2024. "Local Employment Multipliers for Large Publicly Subsidized Firms: Evidence from a Synthetic Control Approach." *Journal of Regional Science* 64 (2): 491–526.

Hendren, Nathaniel, and Ben Sprung-Keyser. 2020. "A Unified Welfare Analysis of Government Policies." *Quarterly Journal of Economics* 135 (3): 1209–1318.

Hershbein, Brad, and Bryan A. Stuart. 2024. "The Evolution of Local Labor Markets After Recessions." *American Economic Journal: Applied Economics* 16 (3): 399–435.

Hoelzlein, Matthias. 2023. "Two-Sided Sorting and Spatial Inequality in Cities." Working Paper.

Hornbeck, Richard, and Enrico Moretti. 2024. "Estimating Who Benefits from Productivity Growth: Local and Distant Effects of City Productivity Growth on Wages, Rents, and Inequality." *Review of Economics and Statistics* 106 (3): 587–607.

Kline, Patrick, and Enrico Moretti. 2014a. "Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority." *Quarterly Journal of Economics* 129 (1): 275–331.

Kline, Patrick, and Enrico Moretti. 2014b. "People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs." *Annual Review of Economics* 6 (1): 629–662.

Marcato, Gianluca, and Lok Man Michelle Tong. 2025. "Supply Constraints and Search Equilibrium in Office Markets." *Journal of Real Estate Finance and Economics* 70 (2): 378–428.

McMillen, Daniel P. 2001. "Nonparametric Employment Subcenter Identification." *Journal of Urban Economics* 50 (3): 448–473.

Mian, Atif, and Amir Sufi. 2014. "What Explains the 2007–2009 Drop in Employment?" *Econometrica* 82 (6): 2197–2223.

Miyauchi, Yuhei, Kentaro Nakajima, and Stephen J. Redding. 2025. "The Economics of Spatial Mobility: Theory and Evidence Using Smartphone Data." *Quarterly Journal of Economics* 140 (4): 2507–2570.

Monte, Ferdinando, Stephen J. Redding, and Esteban Rossi-Hansberg. 2018. "Commuting, Migration, and Local Employment Elasticities." *American Economic Review* 108 (12): 3855–3890.

Moretti, Enrico. 2010. "Local Multipliers." *American Economic Review* 100 (2): 373–377.

Nakamura, Emi, and Jón Steinsson. 2014. "Fiscal Stimulus in a Monetary Union: Evidence from U.S. Regions." *American Economic Review* 104 (3): 753–792.

Neal, Derek. 1995. "Industry-Specific Human Capital: Evidence from Displaced Workers." *Journal of Labor Economics* 13 (4): 653–677.

Neumark, David, and Jed Kolko. 2010. "Do Enterprise Zones Create Jobs? Evidence from California's Enterprise Zone Program." *Journal of Urban Economics* 68 (1): 1–19.

Oh, Ryungha, and Jaeeun Seo. 2023. "What Causes Agglomeration of Services? Theory and Evidence from Seoul." Working Paper.

Patrick, Carlianne. 2016. "Identifying the Local Economic Development Effects of Million Dollar Facilities." *Economic Inquiry* 54 (4): 1737–1762.

Patrick, Carlianne, and Mark Partridge. 2022. "Identifying Agglomeration Spillovers: New Evidence from Large Plant Openings." Working Paper.

Saiz, Albert. 2010. "The Geographic Determinants of Housing Supply." *Quarterly Journal of Economics* 125 (3): 1253–1296.

Santos Silva, J. M. C., and Silvana Tenreyro. 2006. "The Log of Gravity." *Review of Economics and Statistics* 88 (4): 641–658.

Schmidheiny, Kurt, and Sebastian Siegloch. 2023. "On Event Studies and Distributed-Lags in Two-Way Fixed Effects Models: Identification, Equivalence, and Generalization." *Journal of Applied Econometrics* 38 (5): 695–713.

Slattery, Cailin. 2025. "Bidding for Firms: Subsidy Competition in the United States." *Journal of Political Economy* 133 (8): 2563–2614.

Slattery, Cailin, and Owen Zidar. 2020. "Evaluating State and Local Business Incentives." *Journal of Economic Perspectives* 34 (2): 90–118.

Suárez Serrato, Juan Carlos, and Owen Zidar. 2016. "Who Benefits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous Firms." *American Economic Review* 106 (9): 2582–2624.

Suárez Serrato, Juan Carlos, and Owen Zidar. 2023. "Who Benefits from State Corporate Tax Cuts? A Local Labor Market Approach with Heterogeneous Firms: Reply." *American Economic Review* 113 (12): 3401–3410.

Tarczynska, Kasia. 2022. "2022: A Mega-Year for Megadeals!" Good Jobs First. https://goodjobsfirst.org/2022-a-mega-year-for-megadeals/.

Tolbert, Charles M., and Molly Sizer. 1996. "U.S. Commuting Zones and Labor Market Areas: A 1990 Update." ERS Staff Paper AGES-9614, Economic Research Service, U.S. Department of Agriculture, Washington, DC.

Tsivanidis, Nick. 2026. "Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogotá's TransMilenio." *American Economic Review* 116 (2): 418–463.

U.S. Census Bureau. 2025. "Introduction." *LEHD Snapshot Documentation*, Release S2024_R2025Q4. https://lehd.ces.census.gov/data/lehd-snapshot-doc/latest/sections/introduction.html, accessed May 21, 2026.

Veraset. 2022. "Visits." Dewey Data. https://doi.org/10.82551/TP7E-VA02.

Wing, Coady, Seth M. Freedman, and Alex Hollingsworth. 2024. "Stacked Difference-in-Differences." Working Paper 32054, National Bureau of Economic Research.

Yagan, Danny. 2019. "Employment Hysteresis from the Great Recession." *Journal of Political Economy* 127 (5): 2505–2558.

Appendix

Appendix roadmap. Appendix A presents the data construction and additional empirical results. Appendix B derives the model from individual preferences, including employment and location choices, consumption, production, and equilibrium accounts. Appendix C describes quantification and equilibrium computation. Appendix E supports the results in Section 6, and Appendix F supports the counterfactuals in Section 7.

A Data and Additional Empirics

A.1 Award data and event samples

Award records and site verification. The award data begin with three public sources. State economic development agencies publish registers of the awards they approve, Good Jobs First's Subsidy Tracker compiles company-specific awards from state and local disclosures (Tarczynska, 2022), and the Slattery series records the announced packages of large incentive competitions, which often bundle state and local incentives (Slattery, 2025). Where supporting documents exist, I reconcile each recorded amount with agency reports, executed agreements, legislative materials, and local authorizing documents, matching on the recipient, the program, the project location, the award or agreement date, and the contract identifier when one is reported. Because an award names a firm rather than an establishment, I also match each award to its establishment in Data Axle and check the site against a document, preferring primary incentive documents to agency announcements and news reports, and against 2020 tract boundaries and available parcel records. I drop 36 awards that do not attach to a single employer's site, such as grants to prepare industrial land, film credits, and tax increment financing for retail developers, and 11 more whose address cannot be placed in a 2020 tract (Table A.1).

From award records to events. One project often appears in several award records. A state credit, a training grant, and a local abatement may each finance it, and Good Jobs First and a state register may list the same award. I therefore consolidate records into firm-site events. After normalizing firm names, which merges four sites where one firm appears under two spellings, I group the awards to the same firm in the same tract and start a new episode when at least five years separate two awards. Only one site has awards that far apart. When every award at a site promises the same number of jobs, the awards form one event, the jobs count once, and the subsidy is the sum across instruments. When the promises differ, I split the awards by award year and combine the instruments within each year, which turns six sites into nine events. The one exception is a project whose three annual awards report slightly different totals for the same plant, which I keep as a single event. Each event is dated to its first award year, and its promised jobs are the largest promise among its awards. Overlapping valuations of the same award count once, so when a package from the Slattery series already contains a state program's award, I keep the package

figure and do not add the program. These steps turn 299 award records into 226 events involving 212 firms at 201 tracts in 32 states, and 179 of the events rest on a single award.

Realized jobs and payments. Promised jobs come from the award itself, but realized jobs and payments are observed only where a program requires reports and the agency publishes them. Realized employment is observed for 95 events, most of them in Texas, Virginia, North Carolina, and New Jersey, which publish performance reports for their largest programs. The administering agency verified or certified 53 of these counts. The rest are recipients' self-reports republished by the agency, figures carried in the award record, or press reports. When an event has more than one count, I prefer verified or certified counts, then a count attached to a documented post-performance review, then the most recent count, and finally the larger one. The 64 events with a count at least four years after the award, when most of the promised hiring should have occurred, form the descriptive attainment sample in Figure A.1, which is not used in the regressions or the model.

An approved amount is not the amount paid, so I collect payment and clawback records separately where they are public. In Indiana, I match awards to the 2017 Economic Incentives and Compliance Report by recipient, funding program, and exact contract amount, and record the paid or certified amount with its reporting date. In Virginia, I match the Commonwealth's Development Opportunity Fund awards to the state's post-performance reports by grant amount, locality, and performance-agreement date, keep the latest project status among the editions examined, and separate clawback requests from repayments received where the report allows. This yields paid or certified amounts for 22 events and clawback records for five. These records give cumulative payments or recoveries as of a reporting date rather than spending over the completed life of an award.

State and local support. Separating the local share starts with whether a site received any local property-tax relief. Among the 226 events, I establish the local relief for 86, either as a dollar amount (43) or as a documented absence of any abatement (43). For another 13, an abatement is documented but its amount is not published, and the remaining 127 are unresolved. For 33 events, I can document both the state and the local share of the award. Fifteen receive state support only through state programs, fourteen are large packages whose state side I rebuild from the individual programs named in the granting documents, and four have a local deduction whose dollar value I compute from the exempted assessed value. I build the state amount from documented state instruments, rather than as the residual left after subtracting an announced local contribution from the package total. Each component is classified by the revenue source it reduces, so a local property-tax exemption that the state administers remains a local cost. I also set apart spending that does not accrue to the firm, such as public infrastructure, transfers to third parties, and bond face values, and I separate generally available tax preferences from negotiated, firm-specific incentives.

Property-tax relief is researched facility by facility. I first identify the parcel, its owner of record, and the taxing jurisdictions, including the municipality and school district, allowing for

subsidiaries, successor firms, landlords, and industrial development authorities that hold title to the property the recipient occupies. The sources are assessor and auditor records, state abatement registers, industrial development board filings, executed agreements, ordinances, council materials, and project disclosures in Annual Comprehensive Financial Reports. Government-wide abatement totals are not assigned to a firm without project-level evidence. When a source reports an abated assessed value rather than a dollar saving, annual relief equals the jurisdiction's tax rate per dollar of assessed value times the exempted value, and market values are converted only when the assessment basis is known. Under a payment in lieu of taxes, relief is the tax that would have been owed less the required payment. Relief repeated across several state award records counts once, while concessions from separate taxing authorities are added.

Observed annual savings, approval-stage estimates, and projected full-term values are kept separately. When the observation window is incomplete, the full-term value follows the contract's term and schedule, including any recorded cancellation, and otherwise extends the most recent observed annual saving. Cases that require an assumed horizon or a declining deduction schedule are flagged, and all amounts are in 2023 dollars. A search that finds nothing is recorded as missing rather than as zero support, and a documented absence of support is distinguished from an instrument whose amount is unknown. Public-ownership exemptions and substitute cash grants are recorded separately from conventional abatements. Because disclosure and historical coverage vary across jurisdictions, and realized abatements depend on the investment the assessor eventually values, these comparisons describe only the projects whose components can be documented, and their full-term values are estimates rather than verified spending.

Table A.1: From all-state award records to firm-site events

Stage	Change	Count
Raw all-state award records	—	202,070
Package > $1 million or promised jobs > 100	−184,903	17,167
Exactly one usable announced location	−2,587	14,580
Exclude loans	−974	13,606
Append Slattery awards	+385	13,991
Geographic and establishment matching	−5,857	8,134
Award years 2004–2017	−3,503	4,631
Tradable industry	−2,687	1,944
Promised jobs > 500 or > 0.50% of baseline CZ employment	−1,598	346
Restrict to single-employer, site-specific awards	−36	310
Require substantiated address and 2020 tract	−11	299
Consolidate award records into firm-site events	—	**226**

Notes: Counts refer to award records until the final row. CZ employment is measured in the year before the award. Multiple awards for the same firm-site project are combined into one event, dated to the first award year; distinct project commitments and episodes separated by at least five years are kept separate.

The award universe. Table A.2 describes the awards of at least 100 promised jobs in the 29 states with usable records, 2002–2024. Tax credits and rebates are half of the awards but a fifth of the value. Property tax abatements are 7 percent of the awards but a quarter of the value, and they

are local rather than state revenue. The 2 percent of awards that bundle several instruments carry 46 percent of all value, so mean-based statistics are driven by the tail. Most programs require an application before the siting decision, the largest awards are discretionary, and local property-tax deals need a second approval from the county, school board, or industrial development board.

Table A.2: Instrument mix of awards with at least 100 promised jobs, 2002–2024

Instrument	Share of awards	Value ($bn)	Share of value
Tax credit or rebate	52.40%	15.50	21.40%
Grant (incl. loan hybrids)	37.90%	5.70	7.80%
Property tax abatement	6.70%	18.00	24.80%
Megadeal (mixed package)	2.30%	33.20	45.80%
Other (infrastructure, reimbursement, assistance)	0.80%	0.10	0.20%
Total	100%	72.50	100%

Notes: Author's award file, 29 states, 2002–2024, all awards of at least 100 promised jobs. Values are reported package values.

Table A.3 codes the tax base of the 17 programs that carry 99 percent of the credit and rebate dollars. Three fifths are payroll or withholding based, so the firm keeps part of the income tax withheld from its own employees' wages. Only 6 percent are conventional corporate income tax credits. Most "corporate tax credits" in this universe are therefore wage subsidies financed by employees' own withholding. Awards are nominally conditional on job counts, wage floors, and investment minimums, but enforcement ranges from post-performance payment on verified jobs to self-reported counts with weak audit, and reported jobs are usually projections.

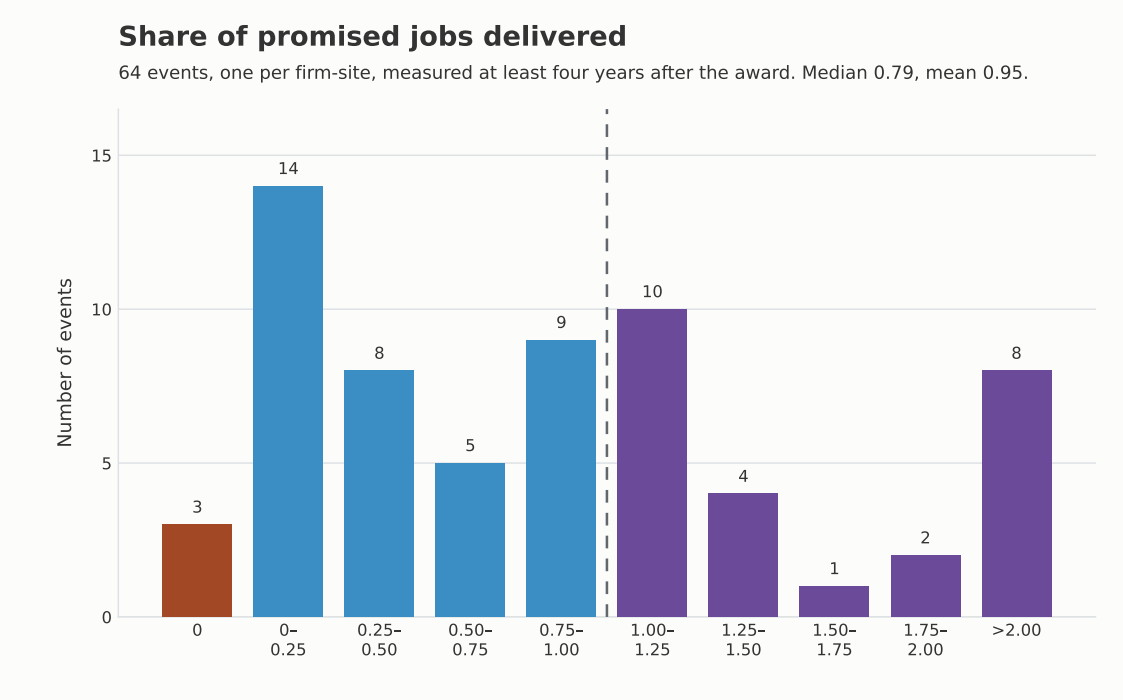
Table A.3: Tax base of credit and rebate awards

Base	Share of credit and rebate value
Payroll or withholding	61.30%
Sales or property, in a state without an income tax	8.10%
Corporate income, conventional	5.90%
Mixed income, sales, and withholding	3.00%
Local or unclassified	3.20%

Notes: Shares of the $15.50 billion coded as tax credit or rebate in Table A.2, after hand-coding the 17 programs that carry 99 percent of it. Two New Jersey programs (25.20 percent of the total) are split between the first and third rows.

Promised and delivered jobs. Figure A.1 uses the 64 firm-site events with reported employment at least four years after award. The median delivered 0.79 of the jobs promised and the mean 0.95. 39 fell short, 17 delivered under a quarter of the promise, 3 recorded zero or negative job gains, and 8 more than doubled the promise. Total realized jobs divided by total promised jobs is 1.06. Because disclosure is selective, these are descriptive statistics rather than estimates of the effect of subsidies.

Figure A.1: Share of promised jobs delivered, firm-site events observed at least four years after award



Notes: 64 events from the consolidated award registry, with employment observed at least four years after award. Each bar counts firm-site events by realized-to-promised employment; the first includes zero and negative gains. The dashed line marks delivery of exactly the promised count. Median 0.79, mean 0.95, pooled ratio 1.06.

The samples used in the paper. The paper draws on several samples of entries, and Table A.4 relates them. The award registry provides the subsidy events. The LEHD entrant sample contains all large tradable openings and expansions in the 29 states with restricted data, whether or not they were subsidized, and supports Facts 1 and 3. The public registry links awards to an entry year, a 2020 tract, and a commuting zone; a candidate enters it only if LODES employment in the entry tract and industry rises by more than 500 jobs relative to baseline in at least one of event years 1 through 5, and when two events occupy the same tract within two years the earlier is kept. The model is quantified for the 55 registry events with complete model inputs. Because the public registry conditions on a realized employment increase, it describes what follows a large entry that occurred and does not identify whether the award caused the firm to locate where it did.

A.2 Public data sources

Employment and commuting. Workplace employment by tract comes from LODES Workplace Area Characteristics, resident workers from Residence Area Characteristics, and tract-to-tract commuting flows from the Origin–Destination files. All tract identifiers are harmonized to 2020 geography. The public service-employment outcome is raw annual workplace employment in NAICS 44–45, 53, 61, 62, 72, and 81, available for 97.48% of tract-years. In the public exposure designs,

Table A.4: Event samples

Sample	Screen	Events	CZs	Used in
Award registry (Table A.1)	Tradable; $>$ 500 jobs or $>$ 0.5% of CZ employment; 2004–2017	226	—	Section 2
with realized jobs $\geq$ 4 years after award	Reported employment	64	—	Figure A.1
LEHD entrant sample	Same size and industry screen; 29 states	400	—	Table 1, Fact 1
pre-entry-residence stack	CZs with one observed event	—	40	Fact 3, Figure 4
new-income stack	CZs with one observed event	—	30	Figure A.8
continuous exposure	All events and tracts	—	100	Figure 5
incumbent-industry event study	Matched control CZs	—	80	Appendix A.4
Public registry	LODES tract-industry employment +500 in years 1–5	124	69	Appendix A.6
boundary comparison	All verified award sites, +5 years	180	86	Section 2, Table A.11
Model sample	Registry events with complete inputs	55	38	Sections 5–7

Notes: Entrant and event counts from restricted data are disclosure-rounded. A dash means the count is not released or not applicable. The LEHD entrant sample includes the subsidized events that meet its screen. The award registry's CZ count is not separately tabulated.

entry tracts and their adjacent tracts receive no exposure, so that exposure is never assigned to employment at the plant itself.

Shopping visits. Shopping flows come from 2019 Veraset smartphone visit records (Veraset, 2022). Explicit home and work records are not counted as visits. Remaining visits to service establishments are aggregated from 2010 block groups to 2020 tracts with the official crosswalk. Each device receives total weight one within a month and industry group, so the annual measure is a device-weighted distribution of visits from each home tract to each destination tract. Visits are counts of trips, not dollars. The visit data cover retail, selected real estate, education excluding colleges, food services, and other services; they contain no health care (NAICS 62) bucket, so that sector enters the employment outcome but not the shopping allocations. Because the visits are observed in 2019, after the events, they describe where service demand is located rather than a predetermined treatment map.

Travel times. Driving times between tract centroids use the 2020 OpenStreetMap road network (OSRM). Same-tract travel is set to five minutes. Gravity comparisons use tract pairs with finite driving time of at most 180 minutes.

Wages. Public workplace wages combine QWI county earnings by education with LODES tract employment counts by education. Because LODES publishes education and industry as separate margins, education categories are allocated to the tradable sector using the tract-year tradable

employment share. This is a proxy for workplace wages by tract, not a joint tabulation.

Property. CoreLogic deed and assessor records over 2004–2020 provide residential and commercial repeat-sales and hedonic price indexes, building square footage, land area, and property counts by 2020 tract. The repeat-sales estimator follows a three-stage Case–Shiller design, and the hedonic regressions include physical characteristics, property type, quality, condition, and missingness indicators. Commercial transactions are sparse in some tracts, and tract-years without support are left missing. Property values are converted to annual rent equivalents in the quantification through the price-to-rent assumptions of Appendix C.

Restricted records. Worker histories and establishment records come from the restricted LEHD and business register (Abowd and Vilhuber, 2011; U.S. Census Bureau, 2025). Entrant hires are observed from 12 quarters before to 12 quarters after hire. A worker's primary employer in a quarter is the employer paying the highest earnings. All restricted statistics in the paper have passed Census Bureau disclosure review.

A.3 Fact 1: Hires, residence, and earnings

Residence and prior workplace. Table A.5 expands Table 2 into ten joint residence–workplace categories for the 551,850 hires in the released residence tabulation, pooling entrant types and CZ sizes. Residence is measured one year before and one year after hire. Let C_e be the event CZ, r_i^- and r_i^{+1} residence before and after hire, and w_i^- the prior workplace. Continuing residents satisfy $r_i^- = r_i^{+1} = C_e$; movers satisfy $r_i^- \neq C_e$ and $r_i^{+1} = C_e$, including those who commuted into the CZ before moving; outside residents satisfy $r_i^{+1} \neq C_e$. Residence and prior workplace carry different information. Continuing residents who previously worked outside the CZ are 9.31 percent of hires, and movers who previously commuted to a job in the CZ are 2.17 percent. Among continuing residents, 9.56 percent have no observed prior job. For these workers I treat no observed job as nonemployment, which includes nonparticipation. For workers who lived outside the CZ before hire, a missing prior job may instead be a job in a state outside the data, so their prior earnings are treated as missing rather than zero.

Market size. Table A.6 splits hires by the released CZ-size class. Continuing residents are 62.27 percent of hires in small CZs, 60.03 percent in medium CZs, and 54.64 percent in large CZs, while the share living outside the CZ after hire rises to 31.26 percent in large markets from about 25 percent in small and medium ones. Movers are 13.09, 15.42, and 14.10 percent of hires, so migration does not rise with market size. These comparisons pool entrant types and do not separate market size from entrant composition.

Prior industry. Table A.7 reports hires' prior industries. Prior jobs in tradables account for 35.41 percent of hires and prior jobs in nontradables for 17.92 percent, of which 10.60 percent of

Table A.5: Residence transitions and prior workplace of entrant hires

Residence before hire	Residence after hire	Prior workplace	Hires	Share of all hires (%)
Outside CZ	Outside CZ	Outside CZ	83,400	15.11
Outside CZ	Outside CZ	Inside CZ	27,700	5.02
Outside CZ	Outside CZ	No observed job	35,600	6.45
Outside CZ	Inside CZ	Outside CZ	37,700	6.83
Outside CZ	Inside CZ	Inside CZ	12,000	2.17
Outside CZ	Inside CZ	No observed job	29,300	5.31
Inside CZ	Outside CZ	Unresolved	7,250	1.31
Inside CZ	Inside CZ	Outside CZ	51,400	9.31
Inside CZ	Inside CZ	Inside CZ	237,000	42.95
Inside CZ	Inside CZ	No observed job	30,500	5.53
Total			551,850	100.00
Among continuing CZ residents (318,900 hires)				
Share with no observed prior job (%)				9.56

Notes: Low- and high-skill entrants and all CZ-size classes are pooled. Residence is measured one year before and one year after hire; workplace refers to the prior job. The ten joint categories are mutually exclusive and exhaustive. Counts are disclosure-rounded, and shares are calculated from those counts, so they may not sum to 100. The final line divides the 30,500 hires with no observed prior job by the 318,900 continuing CZ residents. No observed prior job does not distinguish unemployment from nonparticipation.

Table A.6: Recruitment and residence by CZ size

CZ size	Hires	Continuing CZ residents (%)	Moved into CZ (%)	Lives outside CZ after hire (%)
Small	106,150	62.27	13.09	24.63
Medium	171,900	60.03	15.42	24.55
Large	273,800	54.64	14.10	31.26

Notes: Entrant types are pooled within each CZ-size class. Each percentage uses all hires in its row as the denominator. Continuing residents live in the event CZ one year before and one year after hire; movers live outside it before and inside it after; the final category lives outside it after hire. Counts are disclosure-rounded. Small, medium, and large are the released size labels; population cutoffs are not released.

all hires came from nontradable employers inside the event CZ. A further 17.45 percent have no observed prior job.

Earnings. For released group g with n_g hires, let q_g^+ and q_g^- be the after-hire and prior annual earnings quasi-medians, and let $\mathcal{O}(\mathcal{G})$ be the groups in a set $\mathcal{G}$ with observed prior earnings. The

Table A.7: Prior industry and workplace of entrant hires

Prior industry / job status	All hires (%)	Prior job inside CZ (%)	Prior job outside CZ (%)
Low-skill nontradables	13.06	7.32	5.74
High-skill nontradables	4.87	3.28	1.59
Low-skill tradables	15.42	10.40	5.01
High-skill tradables	19.99	12.78	7.21
Other observed industries	29.22	17.31	11.91
No observed prior job	17.45	—	—

Notes: Entrant types and worker skills are pooled. Each percentage uses all 550,700 hires in the prior-industry tabulation as its denominator; this sample is separate from the 551,850 hires in the residence tabulation. Low and high skill refer to the prior employer's industry. The inside- and outside-CZ columns split each observed-industry row geographically. Counts are disclosure-rounded.

reported after-hire earnings and earnings increase are

$$Q_{\mathcal{G}}^{+} = \frac{\sum_{g\in\mathcal{G}} n_g q_g^{+}}{\sum_{g\in\mathcal{G}} n_g}, \qquad D_{\mathcal{G}} = \frac{\sum_{g\in\mathcal{O}(\mathcal{G})} n_g (q_g^{+} - q_g^{-})}{\sum_{g\in\mathcal{O}(\mathcal{G})} n_g}. \tag{A.1}$$

These are count-weighted averages of group quasi-medians, not medians of individual earnings changes. Among the 454,600 hires with observed prior earnings, after-hire earnings average $85,500 and the increase $22,500. The increase is $18,400 for hires whose prior job was in the CZ and $29,300 for hires whose prior job was outside it (Table A.8). The 96,100 hires with no observed prior job earn $72,200 after hire and have no reported increase. The increases can reflect hours as well as pay rates.

Table A.8: Earnings after hire and observed earnings increases

Hire group	Hires	2023 $ thousands After hire	Increase
All hires with observed prior earnings	454,600	85.5	22.5
Prior workplace inside CZ	281,350	84.9	18.4
Prior workplace outside CZ	173,250	86.4	29.3
Low-skill workers, observed prior earnings	299,600	61.7	18.3
High-skill workers, observed prior earnings	155,000	131.4	30.7
No observed prior job	96,100	72.2	—

Notes: Entrant types are pooled. Earnings are annual, in thousands of 2023 dollars, and are count-weighted averages of released group quasi-medians; increases are the corresponding averages of after-minus-before quasi-medians. The prior-workplace rows and the worker-skill rows are alternative splits of the same 454,600 hires. Missing prior earnings are not treated as zero. Counts are disclosure-rounded; the earnings sample contains 550,700 hires.

Payroll by residence. Table A.9 splits the payroll shares in Table 2 by worker skill. Of the $69.035 billion in entrant payroll, 71.69 percent is paid to workers living inside the event CZ one year after hire. The inside-CZ total combines continuing residents and movers.

Table A.9: Entrant payroll by post-hire residence and worker skill

	Annual payroll (2023 $ billions)			
Residence one year after hire	Low-skill workers	High-skill workers	Total payroll	Share of total payroll (%)
Inside the CZ	23.852	25.639	49.491	71.69
Outside the CZ	9.284	10.260	19.544	28.31
Total	**33.136**	**35.899**	**69.035**	**100.00**

Notes: Low and high skill refer to individual workers; entrant types are pooled. Residence is measured one year after hire. Payroll is annual, in billions of 2023 dollars, from released earnings sums. These are gross payments by post-hire residence, not additional resident income.

The resident-income accounting. Fact 1 allocates each $100 of entrant payroll among workers living outside the CZ, the prior earnings of continuing residents, and new resident income. This paragraph states the object being approximated and the assumptions the approximation uses. For entrant e, let H be its hire cohort and $P^H_{e,0}$ the entrant's payroll to that cohort in a reference year. For horizon h, let $\mathcal{W}^-_h$ and $\mathcal{W}^+_h$ be annual earnings windows before and after hire, r^-_h and r^+_h the residence reference quarters, $Y^{\text{cov}}_{i,\mathcal{W}}$ worker i's covered earnings from all employers in window $\mathcal{W}$, and $Y^e_{i,\mathcal{W}}$ earnings paid by the entrant. The change in resident earnings and the entrant payroll paid to outside residents are

$$\Delta Y^{H,R}_{e,h} \equiv \sum_{i \in H} \left[\mathbf{1}\{R_{i,r^+_h} = C_e\} Y^{\text{cov}}_{i,\mathcal{W}^+_h} - \mathbf{1}\{R_{i,r^-_h} = C_e\} Y^{\text{cov}}_{i,\mathcal{W}^-_h} \right], \tag{A.2}$$

$$P^{H,C}_{e,h} \equiv \sum_{i \in H} \mathbf{1}\{R_{i,r^+_h} \neq C_e\}\, Y^e_{i,\mathcal{W}^+_h}, \tag{A.3}$$

and the resident-earnings ratio is $\iota^{\text{earn}}_{e,h} \equiv \Delta Y^{H,R}_{e,h} / P^H_{e,0}$. Equation (A.2) requires earnings classified jointly by residence before and after hire. The released tabulations classify residence (Table A.5) and earnings (Table A.8) separately, so the $100 allocation in Fact 1 approximates it under three assumptions. First, each residence–workplace category's earnings equal the released quasi-median for its prior workplace, pooling prior industries and worker skills within entrant type. Second, continuing residents with no observed prior job had zero prior earnings. Third, prior and after-hire allocations are each scaled to match the released payroll totals by entrant type and post-hire residence. The approximation omits replacement hiring at the employers the hires left and any wage and employment response elsewhere in the CZ, which is why Fact 1 presents its $32 as a lower bound on the resident income an entry creates.

No double counting. The accounting above enters the model exactly once. In Section 4, entrant jobs enter tradable labor demand, the resulting wages accrue to the workers who fill them, and their earnings pass once through the housing and service shares and the 70/30 home/work shopping split, with inbound commuters' home-origin spending and outbound commuters' workplace-origin spending leaving the local account. Adding the Fact 1 allocation as a second spending shock would count the same payroll twice. The model's recruitment of entrant workers is an equilibrium outcome rather than an imposed share.

Education and reasons for moving. The higher migrant share at high-skill entrants in Section 2 is consistent with public evidence on why workers move. The CPS ASEC for 2007–2019 records each respondent's main reason for moving. For respondents aged 25–64, I estimate survey-weighted linear probability models of five mobility outcomes on an indicator for a bachelor's degree, controlling for survey year, five-year age group, sex, race and ethnicity, and marital status, and report average standardized predictions by education group with confidence intervals from the 160 CPS replicate weights (Table A.10). College graduates are no more likely to move in a given year (11.44 versus 11.35 percent), but they are 2.07 times as likely to move for a new job or job transfer (1.87 versus 0.90 percent), and among movers 16.41 percent of graduates cite a new job against 8.01 percent of others. Education is associated not with whether people move but with whether the move is tied to a job.

Table A.10: Educational attainment and reasons for residential mobility

Outcome	Below BA (%)	BA+ (%)	Difference pp [95% CI]	Rate ratio [95% CI]
Any residential move	11.35	11.44	0.10 [−0.07, 0.26]	1.01 [0.99, 1.02]
Employment-related move	2.05	3.24	1.18 [1.09, 1.28]	1.58 [1.53, 1.63]
New job or job transfer	0.90	1.87	0.97 [0.90, 1.03]	2.07 [1.98, 2.17]
Employment reason, among movers	18.20	28.38	10.18 [9.41, 10.94]	1.56 [1.51, 1.61]
New job or transfer, among movers	8.01	16.41	8.40 [7.82, 8.98]	2.05 [1.96, 2.14]

Notes: Public-use CPS ASEC, 2007–2019, respondents aged 25–64. BA+ denotes a bachelor's degree or higher; associate degrees are in the below-BA group. The first three rows use all respondents as the denominator; the last two condition on having moved. Rates are average standardized predictions from survey-weighted linear probability models controlling for year, age group, sex, race and ethnicity, and marital status. Confidence intervals use CPS successive-difference replicate weights.

A.4 Incumbent adjustment

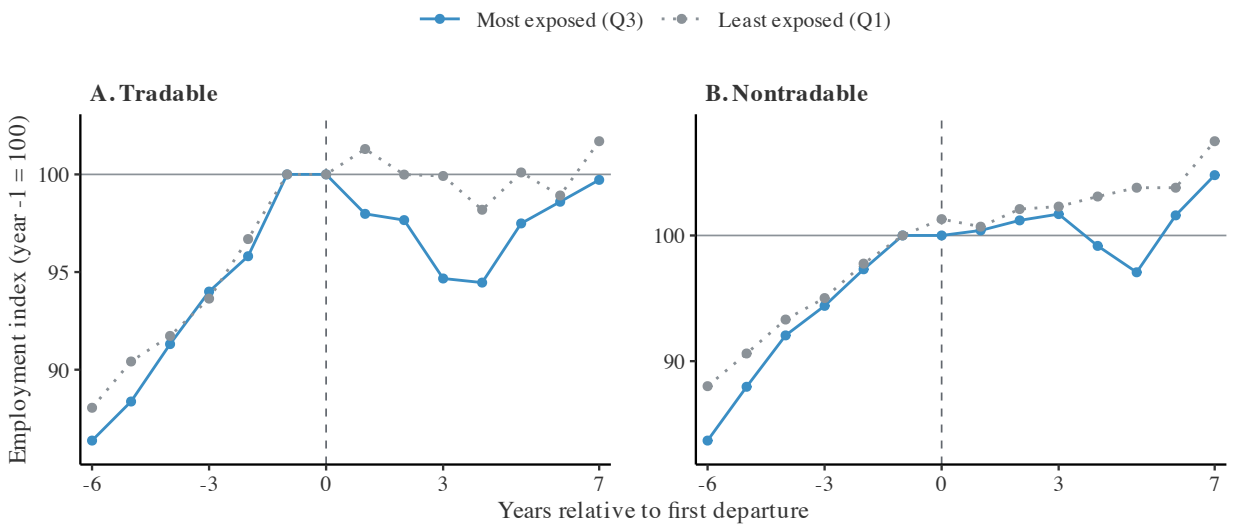
Source establishments. For each hire who previously worked in the CZ, the source establishment is the hire's most recent main job. Exposure is the share of a source establishment's payroll that the entrant hired away, measured the year before its first worker left. Within each size bin, the figures compare the most exposed third of source establishments with the least exposed third.[32]

[32]Exposure terciles are formed separately within four employment size bins: 6–19, 20–49, 50–249, and 250 or more workers.

Each index equals 100 in the year before the first departure. The all-establishment employment index counts establishments that later close as zero, while the active-establishment employment and pay indices follow only establishments that remain open.

Among tradable establishments that remain open, employment at the most exposed falls about 5.3 jobs per 100 baseline jobs behind the least exposed establishments three years after the first departure, and the gap narrows to about 2 jobs per 100 by year seven (Figure A.2). The year-seven gap is about 2.7 jobs per 100 for nontradable establishments.[33] Because every establishment in the sample exists at the first departure, the all-establishment index rises before that year and falls after it for both exposure groups, so the comparison that matters is between the two groups at the same point in time. Once closures are counted, the most exposed tradable establishments end seven years after the first departure about 13.5 jobs per 100 baseline jobs below the least exposed ones (Figure A.3), and those still operating raised average pay by 14 percent, compared with 15 percent at the least exposed establishments (Figure A.4). These are descriptive paths without confidence intervals.

Figure A.2: Employment at active source establishments by exposure

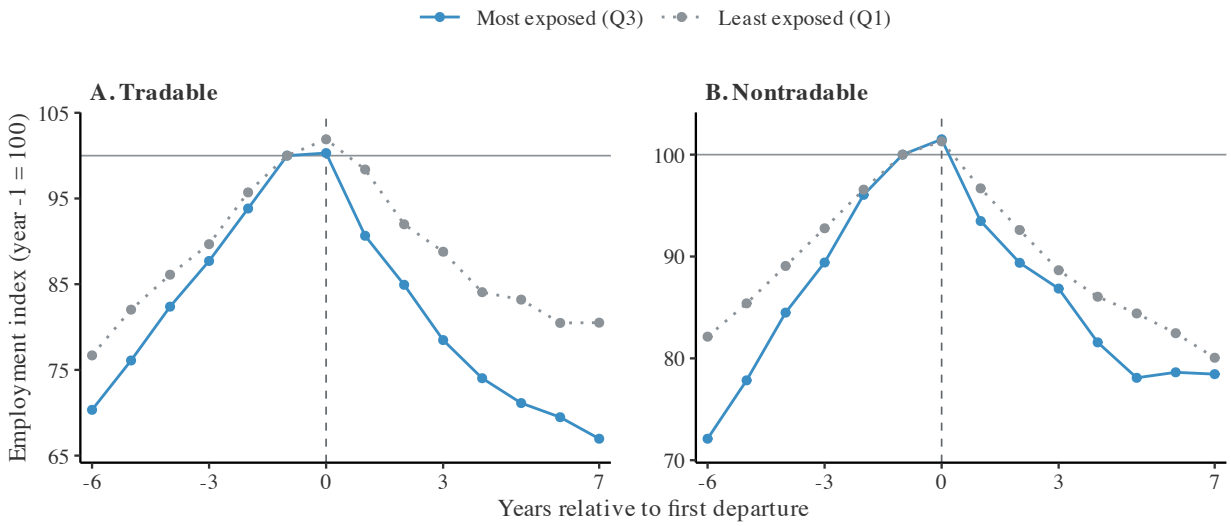


Notes: Released employment index for source establishments that remain active, for the most-exposed (Q3) and least-exposed (Q1) thirds of source establishments, by tradable and nontradable sector. Exposure is measured as the share of an establishment's baseline payroll lost to the entrant. Year 0 is the first departure; all indices equal 100 in year −1. Missing values at years −1 and 0 are displayed at 100. Each panel has its own scale. Figures A.3 and A.4 show employment across all establishments and average pay at active establishments.

Labor-market relatedness and incumbent employment. To investigate whether large tradable entrants displace employment among incumbent firms through competition for workers, I examine whether incumbent industries that are more closely connected to the entrant in the labor

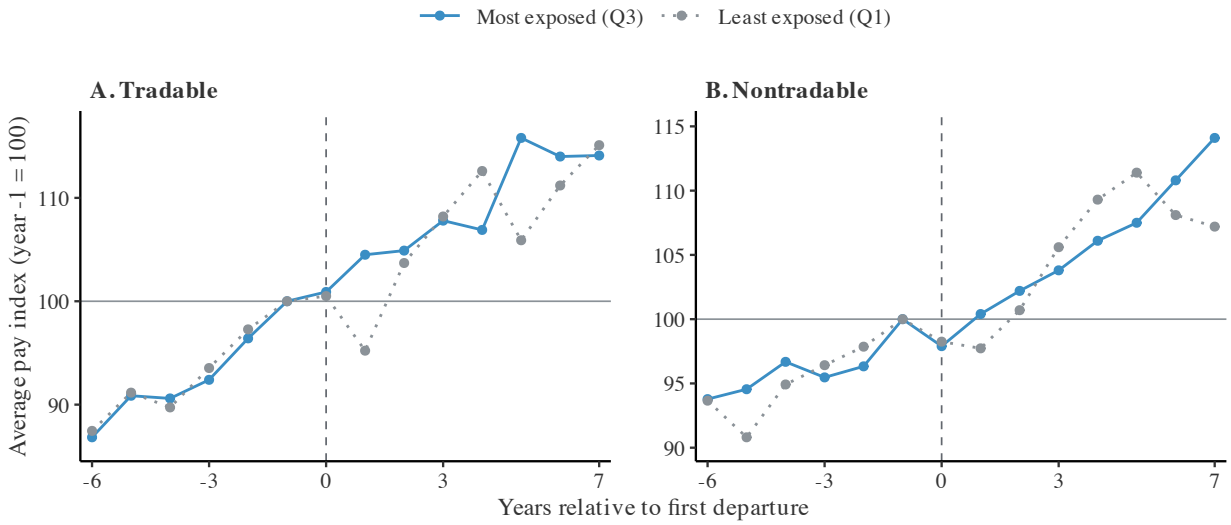
[33]Sectors are defined by the source establishment's four-digit NAICS industry. Tradable industries follow the classification of Mian and Sufi (2014), excluding anything in NAICS 61, 72, and 81 and including all of manufacturing (NAICS 31–33). Nontradable industries are the local services defined in Section 2.

Figure A.3: Employment at all source establishments by exposure



Notes: Released employment index for all source establishments, which counts establishments that later close as zero, for the most-exposed (Q3) and least-exposed (Q1) thirds of source establishments, by tradable and nontradable sector. Groups and display conventions follow Figure A.2.

Figure A.4: Average pay at active source establishments by exposure



Notes: Released average-pay index for source establishments that remain active, for the same groups and sectors. Groups and display conventions follow Figure A.2.

market experience larger employment declines. I construct a predetermined measure of labor-market relatedness using national worker flows I construct from the LEHD between three-digit NAICS industries. Industries that historically supply a larger share of workers to the entrant's industry are expected to face greater recruitment competition following entry. For each entry event

e and incumbent industry k, I define labor-market exposure as

$$L_{ek} = \frac{J_e \, s_{k \to h(e)}}{E^0_{ek}}, \tag{A.4}$$

where J_e denotes entrant employment, $s_{k \to h(e)}$ is the predetermined national worker-flow relatedness between incumbent industry k and the entrant's industry $h(e)$, and E^0_{ek} is baseline incumbent employment in industry k in the treated commuting zone. I normalize this measure by its employment-weighted interquartile range, so the coefficients correspond to a one-IQR increase in labor-market exposure. I estimate the event-study specification

$$y_{eckt} = \alpha_{eck} + \delta_{ect} + \gamma_{ekt} + \sum_{\tau \neq -2} \beta_\tau \, \mathbf{1}\{t - T_e = \tau\} \, D_{ec} \, \widetilde{L}_{ek} + \varepsilon_{eckt}, \tag{A.5}$$

where y_{eckt} is log incumbent employment, or log average pay in industry k, commuting zone c, and year t, for event e. The indicator D_{ec} identifies the treated commuting zone, and $\widetilde{L}_{ek}$ is normalized labor-market exposure. The fixed effects absorb time-invariant event–CZ–industry differences, event–CZ–year shocks, and event–industry–year shocks. The omitted period is two years before entry.

For each large-firm entry event, I select five comparison commuting zones using propensity-score matching. Candidate commuting zones are restricted to the same Census division and must not experience another identified large-firm entry within five years of the event. Matching covariates, measured two years before entry, include baseline employment, number of establishments, population, demographic composition, residential rents, and the entrant industry's local employment concentration.

The sample follows a fixed cohort of incumbent establishments identified before entry, retaining establishments that subsequently close with zero employment. Regressions are weighted by baseline incumbent-industry employment, with matched control locations collectively receiving the same weight as the treated location for each event-industry cell. Standard errors are clustered by event commuting zone and incumbent industry. The coefficients β_τ capture differential changes in incumbent outcomes associated with greater labor-market exposure in treated locations relative to matched controls. I summarize the post-entry response using the average coefficient over years two through five and examine the joint significance of pre-entry and post-entry coefficients.

Figure 2 plots the labor coefficients from the joint specification. Tradable industries more connected to the entrant through worker flows lose employment in years four through six, by at most 0.06 percent per interquartile range of exposure, and the year-seven interval includes zero. Average pay in those tradable industries is also slightly lower after entry, by about 0.03 percent per interquartile range, while nontradable industries show no clear change. Several pre-entry intervals also exclude zero, so these paths are suggestive rather than clean event studies.

A.5 Fact 2: Shopping geography and openness

Service expenditure shares. The 2024 BEA PCE Bridge maps each NIPA consumption category to input–output commodities and reports producers' values, transportation costs, and wholesale and retail margins. I define service expenditure N as household consumption of educational services (NAICS 61), health care and social assistance (62), accommodation and food services (72), and other services (81), each at producers' value, plus the retail margin on goods sold by retailers (44–45). Using the margin rather than the full purchase price keeps the manufactured value of goods out of retail. This gives $N = \$9.04$ trillion in 2024, or $N/Y^d = 0.41$ of disposable personal income Y^d, which was \$21.92 trillion (NIPA Table 2.1). Housing $H = \$3.15$ trillion is the PCE housing-services category, including imputed rent on owner-occupied housing; a \$3.29 billion overlap with group housing is removed. The share of after-housing income spent on services is

$$s_N^{AH} = \frac{N - 0.00329}{Y^d - H} = \frac{9.04}{21.92 - 3.15} = 0.48,$$

or 49.34% if utilities are also treated as housing. These are average national shares of total household consumption, including services purchased on households' behalf by employers and governments, not marginal propensities, and the sectors are locally oriented rather than perfectly nontraded.

Openness at county and CZ boundaries. Figure A.5 compares, for each county and CZ, the share of residents' primary jobs located inside the area with the share of residents' service visits that end inside it. The commuting share uses 2019 LODES resident workers; the visit share uses raw 2019 Veraset visits in retail, selected real estate, education excluding colleges, food services, and other services. Neither is a dollar share.

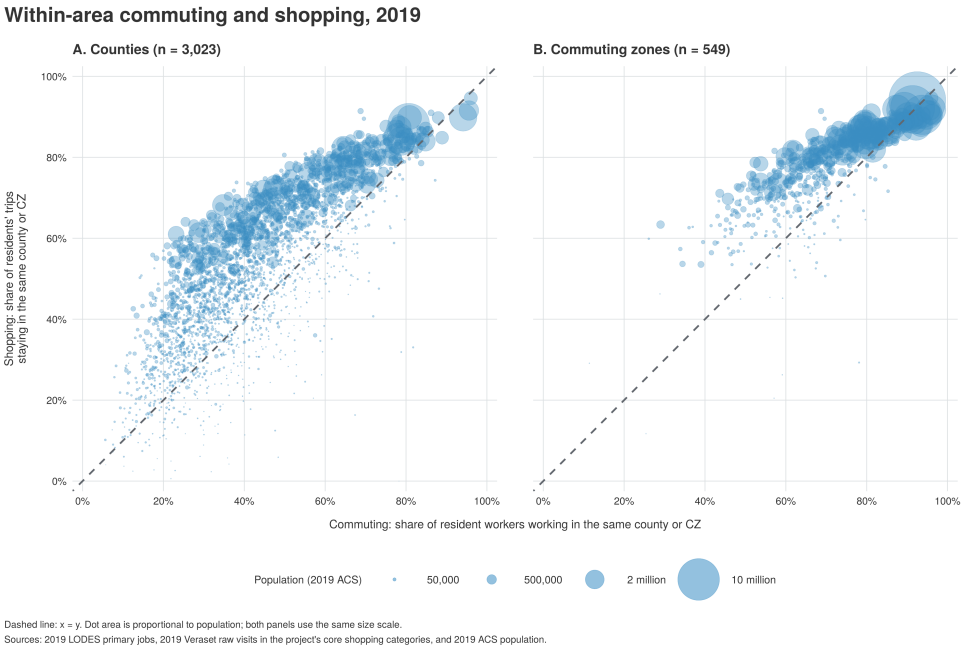
The entrant-level comparison in Section 2 uses the 180 verified award episodes at 167 workplace tracts in 86 CZs, with 4 episodes more than 100 km from their CZ centroid excluded. Let F_{ie,t_e+5} be the number of primary-job workers living in tract i and employed in the event tract e five years after the award. The share living inside boundary B is

$$S_e(B) = \frac{\sum_{i \in B} F_{ie,t_e+5}}{\sum_i F_{ie,t_e+5}}, \tag{A.6}$$

for the event county, the CZ, and the tracts within 60 driving minutes of the event tract. These are post-award levels for all workers in the event tract, not the entrant's own workforce. Entrant industries are classified as high skill when their four-digit industry's college share in the 2000 Census is at or above the national median of 20.66%, giving 132 high-skill and 46 low-skill episodes. To separate accessible population from administrative geography, I also compute the 2019 population within 40 driving minutes of the event tract,

$$P_e^{40} = \sum_i \text{Pop}_{i,2019}\ \mathbb{1}\{\tau_{ie} \leq 40 \text{ minutes}\}, \tag{A.7}$$

Figure A.5: Within-area commuting and shopping, 2019



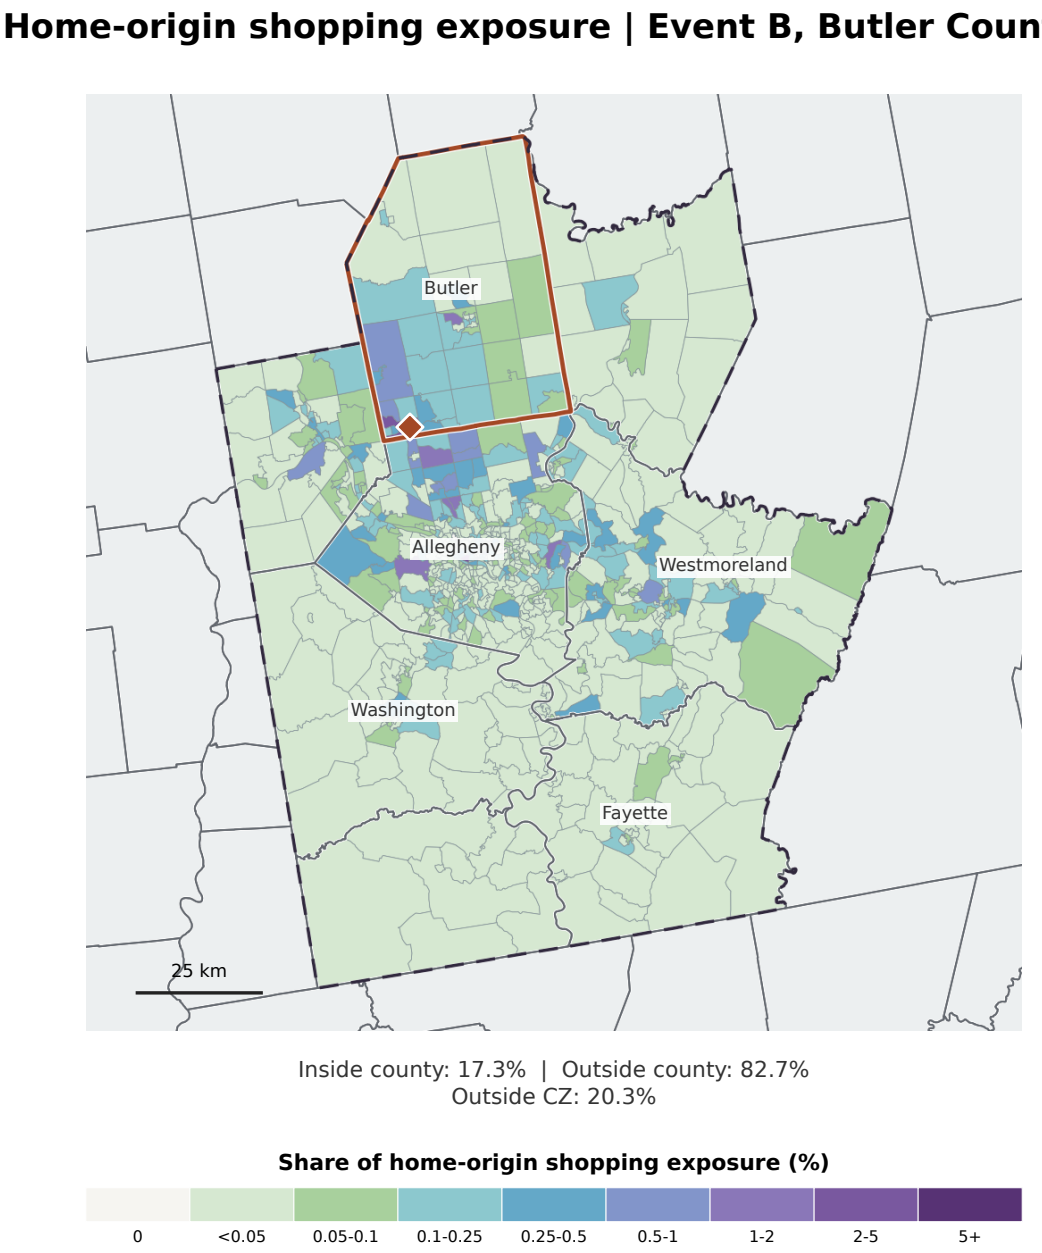
Notes: Each dot is a county (3,023) or CZ (549). The horizontal axis is the share of resident primary-job workers in LODES 2019 who work inside the area; the vertical axis is the share of residents' raw Veraset visits in NAICS 44, 45, selected 53, 61 (excluding 6113), 722, and 81 that end inside it. Cross-county flows within a CZ count as internal. Dot area is proportional to 2019 population; the line is equality.

and split episodes at common terciles (448,246 and 1,735,422 residents). Table A.11 reports the shares. In the lower and middle terciles, high-skill sites keep a larger share of their workers inside the county than low-skill sites, the reverse of the unconditional comparison, so the aggregate skill contrast partly reflects where each type of entrant locates. The upper tercile has only 2 low-skill events.

Figure A.6 maps where the workers at the entrant's tract in Figure 1 shop five years after entry, assigning each worker the 2019 visit pattern of residents of their home tract.

Where within-CZ visits go. Table 3 in Section 3 splits within-CZ visits into three mutually exclusive destinations. A same-tract visit ends in the device's home tract; an adjacent-tract visit ends in a different 2020 tract sharing a boundary point with the home tract; a nonlocal hub visit ends in a shopping-hub footprint outside the home tract and its neighbors. For the core sectors (NAICS 44–45, 722, 81), 10.60 percent of within-CZ visits stay in the home tract, 23.81 percent go to an adjacent tract, and 29.25 percent to a nonlocal hub. Hubs of all kinds receive 43.40% of within-CZ visits; 14.10% of within-CZ visits are to a hub that is also in the home tract or an adjacent tract, so 67.50% of hub visits are away from home. Across all destinations, 84.80% of core-sector visits stay in the home CZ. The core combines NAICS 44, 45, selected NAICS 53, restricted NAICS 61, NAICS 722, and NAICS 81. Selected NAICS 53 consists of 5311 (Lessors of Real Estate), 5312 (Offices of Real Estate Agents and Brokers), 5313 (Activities Related to Real

Figure A.6: Where the workers at the entrant's tract shop



Notes: Share of the shopping visits of the workers at the entrant's tract five years after entry that end in each destination tract, in percent. Each worker is assigned the 2019 visit pattern of residents of their home tract across retail (NAICS 44–45), food services (722), other services (81), and non-college education (61), weighted by public LODES counts of primary jobs in the broad All Other Services group (SI03) at the tract by residence. Workers' own purchases are not observed. Shares under the map include destinations outside the CZ. The diamond marks the entrant's tract and the outlined area its host county.

Table A.11: Share of workers living inside the boundary, by nearby population and entrant industry

Entrant industry	Events	County	CZ	60-minute area
		Mean share in percent (SD in percentage points)		
All retained events				
High-skill	132	44.60 (18.50)	78.50 (11.40)	83.50 (7.50)
Low-skill	46	48.50 (15.20)	75.10 (12.20)	77.70 (9.90)
Lower tercile: at most 448,246 residents				
High-skill	25	57.60 (12.90)	79.60 (8.80)	81.40 (7.70)
Low-skill	34	50.50 (15.70)	77.00 (10.10)	76.10 (10.10)
Middle tercile: 448,246–1,735,422 residents				
High-skill	49	45.60 (19.00)	81.20 (9.50)	81.80 (8.50)
Low-skill	10	42.20 (13.00)	70.80 (15.20)	81.00 (8.60)
Upper tercile: more than 1,735,422 residents				
High-skill	58	38.30 (17.20)	75.70 (13.20)	85.90 (5.70)
Low-skill	2	45.90 (16.70)	63.70 (25.50)	87.80 (4.20)

Notes: Each event receives equal weight. The denominator is all LODES primary jobs in the event workplace tract in award year +5, not the entrant workforce. Columns change only the residential boundary. Parentheses contain cross-event standard deviations, not standard errors.
Nearby population is 2019 ACS population summed across 2020 tracts whose reference points are at most 40 driving minutes *to* the event tract, irrespective of county or CZ membership. Whole-tract populations are included by reference-point membership. This is an accessible-population measure, not population density or a pre-entry covariate.
Common tercile cutoffs use all 180 retained award episodes, including 2 with unresolved industry skill classifications. Those two episodes are omitted from the displayed high/low rows. Ties stay together; displayed cutoffs are rounded to the nearest resident. The upper-tercile low-skill group has only 2 events.
High skill is at or above the nationwide median four-digit-industry college share in Census 2000; pooled industry matches are flagged in the source data. See Appendix A.5 for the sample and coverage checks.

Estate), 5321 (Automotive Equipment Rental and Leasing), 5322 (Consumer Goods Rental), and 5323 (General Rental Centers). Restricted NAICS 61 includes Educational Services other than 6113 (Colleges, Universities, and Professional Schools). Mobility records for NAICS 62 (Health Care and Social Assistance), including ambulatory health-care providers, hospitals, nursing and residential-care facilities, and social-assistance providers, are excluded for privacy reasons. NAICS 721 is not included in the core or reported separately. Table A.12 reports the cumulative driving-time distribution of all visits, including those that cross CZ boundaries.

Shopping hubs. Hubs are built from destination activity, adapting the employment-subcenter method of Giuliano and Small (1991) and McMillen (2001) to visits. Let V_j be all 2019 device-weighted core-sector visits received by tract j, and $\widetilde{V}_j = V_j + \sum_{k \in \mathcal{N}(j),\, c(k)=c(j)} V_k$ its sum with its same-CZ neighbors. A tract is a candidate peak when $\widetilde{V}_j$ is at least as large as that of every adjacent tract in its CZ. Peaks are processed from strongest to weakest, and a weaker peak is

Table A.12: Cumulative driving time of U.S. service visits, 2019

Sector	≤ 5	≤ 10	≤ 15	≤ 20	≤ 30	≤ 40	≤ 60	> 60
NAICS 44	18.99	38.11	52.51	62.61	74.86	81.16	86.82	13.18
NAICS 45	14.86	34.60	50.95	62.16	75.23	81.72	87.43	12.57
NAICS 722	16.25	34.75	49.25	59.60	72.23	78.79	84.80	15.20
NAICS 81	19.20	36.28	50.92	61.99	75.72	82.59	88.45	11.55
Core	*16.86*	*35.52*	*50.48*	*61.11*	*73.97*	*80.52*	*86.38*	*13.62*
Selected NAICS 53	9.26	22.44	35.35	45.79	59.50	66.84	73.76	26.24
Restricted NAICS 61	11.28	26.24	41.35	53.74	69.77	78.22	85.99	14.01

Notes: Percentages of annual device-weighted visits, including visits that cross CZ boundaries. Columns are cumulative minutes of tract-center driving time on the 2020 road network; same-tract visits are assigned five minutes. Pairs more than 100 straight-line miles apart, routed times above 60 minutes, and unreachable pairs are in the > 60 column. The core combines NAICS 44, 45, selected NAICS 53, restricted NAICS 61, NAICS 722, and NAICS 81. Selected NAICS 53 consists of 5311 (Lessors of Real Estate), 5312 (Offices of Real Estate Agents and Brokers), 5313 (Activities Related to Real Estate), 5321 (Automotive Equipment Rental and Leasing), 5322 (Consumer Goods Rental), and 5323 (General Rental Centers). Restricted NAICS 61 includes Educational Services other than 6113 (Colleges, Universities, and Professional Schools). Mobility records for NAICS 62 (Health Care and Social Assistance), including ambulatory health-care providers, hospitals, nursing and residential-care facilities, and social-assistance providers, are excluded for privacy reasons. NAICS 721 is not included in the core or reported separately.

merged into a stronger one when the mean driving time between them is at most ten minutes. A hub's footprint $\mathcal{F}_g$ is the union of its peak tracts and their adjacent tracts, with shared tracts assigned to the stronger hub, and its share of destination activity is $H_g = \sum_{j \in \mathcal{F}_g} V_j / \sum_{j:c(j)=c(g)} V_j$. A hub qualifies under the 2 percent definition when $H_g \geq 0.02$; the 1 and 5 percent definitions use 0.01 and 0.05. The same footprints are used for every sector in Table 3.

Hubs and market size. Across the 446 CZs with 2019 population of at least 50,000, the number of hubs rises with population. A log-link count model of the hub count on log population gives an elasticity of 0.59 under the 1 percent definition, 0.43 under the 2 percent definition, and 0.10 under the 5 percent definition, which mechanically caps how many hubs can qualify (Table A.13).

Table A.13: Scaling of within-CZ shopping hubs with CZ population

Minimum hub share	Preferred model	Intercept	Robust SE	Elasticity	Robust SE	95% CI
1%	Negative binomial	-5.95	(0.30)	0.59	(0.02)	[0.54, 0.63]
2%	Negative binomial	-3.99	(0.32)	0.43	(0.03)	[0.38, 0.48]
5%	Poisson	-0.20	(0.30)	0.10	(0.02)	[0.06, 0.15]

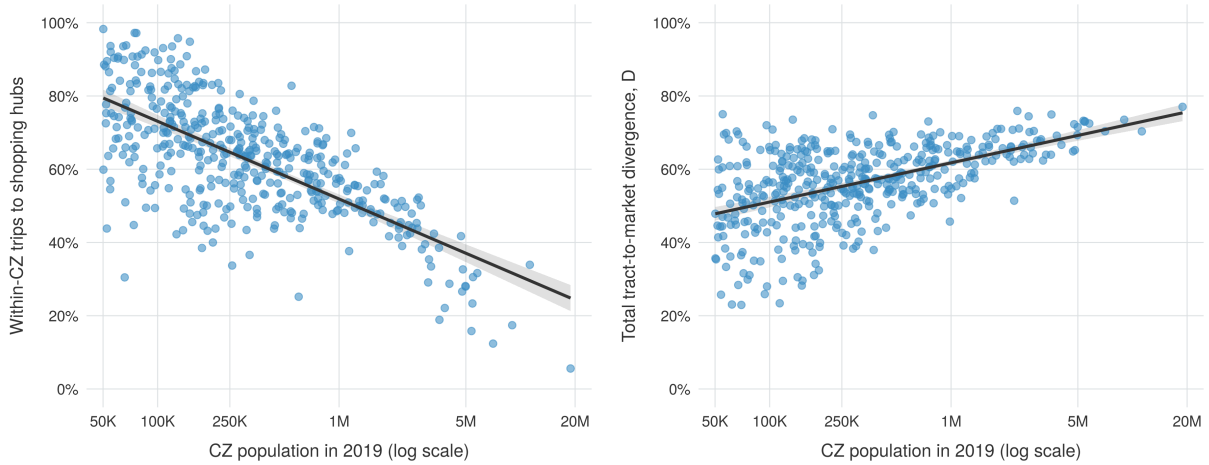
Notes: Each row reports the valid AIC-preferred count model for one hub-mass threshold. The sample contains 446 CZs with 2019 population of at least 50,000. The dependent variable is the number of within-CZ hubs. Standard errors are heteroskedasticity-robust HC1 estimates. At the 5% threshold, the negative-binomial dispersion collapses to the Poisson boundary and is excluded from model comparison.

A tract's shopping versus its market's. Fact 3 assigns exposure to the home tract and its neighbors rather than through the observed shopping network, on the grounds that the home-and-adjacent component is the part of shopping that differs across tracts. This paragraph measures that. For home tract i in CZ c, let p_{ij} be the share of its device-weighted within-CZ visits ending in tract j, and let

$$q_{-i,j} = \frac{\sum_{k \in c,\, k \neq i} v_{kj}}{\sum_{k \in c,\, k \neq i} \sum_{\ell \in c} v_{k\ell}} \tag{A.8}$$

be the destination shares of all other residents of the CZ. The total-variation distance $D_i = \sum_j \max\{p_{ij} - q_{-i,j}, 0\}$ is the share of tract i's visits that would have to be reassigned to make its pattern match its market's. Across 438 CZs with population of at least 50,000, 34.40% of within-CZ visits end in the home tract or an adjacent tract, against 3.20% under the other residents' pattern, and the home-and-adjacent set accounts for 48 percent of the pooled positive deviation $\sum_i V_i D_i$ (weighting origins by their visits V_i). Figure A.7 shows how the two parts of this vary with market size. The share of within-CZ visits going to 2 percent hubs falls from three quarters in the smallest tract-count decile to two fifths in the largest (correlation -0.70 with log population), while the visit-weighted mean of D_i rises from about one half to two thirds (correlation 0.52). Small markets send most visits to a few shared hubs, so tracts look like their markets; large markets spread visits over many destinations, and tracts differentiate.

Figure A.7: Shopping hubs and total tract-to-market divergence by market size



Notes: Each dot is one of the 438 CZs with 2019 population of at least 50,000 and positive within-CZ shopping support. Left: share of residents' within-CZ device-weighted visits received by 2 percent hub footprints. Right: visit-weighted mean of D_i across the CZ's home tracts. Lines are OLS fits in log population with pointwise 95 percent HC1 bands.

Shopping and commuting gravity. This paragraph documents the comparison of travel-time decay behind Figure 3 and equation (1). It is distinct from the sector-by-CZ shopping gravity used to quantify the model in Section 5, which adds same-tract and same-county terms. For shopping, F^S_{ijc} is the number of 2019 device-weighted visits from residents of i to establishments in j, pooling

NAICS 44–45, 5311–5313, 5321–5323, 61 (excluding 6113), 722, and 81. For commuting, F^C_{ijc} is the 2019 LODES count of workers living in i and working in j. For each mode m and CZ c, Poisson pseudo-maximum likelihood solves

$$(\widehat{\alpha}^m_c, \widehat{\delta}^m_c, \widehat{\beta}^m_c) = \arg\max_{\alpha,\delta,\beta} \sum_{(i,j)\in\mathcal{J}_c} \left[F^m_{ijc} \log \mu^m_{ijc} - \mu^m_{ijc}\right], \qquad \mu^m_{ijc} = \exp\left(\alpha^m_{ic} + \delta^m_{jc} + \beta^m_c t_{ijc}\right), \quad \text{(A.9)}$$

with standard errors two-way clustered by origin and destination tract. PPML accommodates the many tract pairs with zero flows, which carry information about decay, and avoids the retransformation problems of log-linear OLS (Santos Silva and Tenreyro, 2006). Both modes use the same tract-pair support: all non-water pairs within a CZ with finite driving time of at most 180 minutes, with same-tract pairs at five minutes and absent flows entered as zero. A CZ needs at least ten candidate pairs and ten positive-flow pairs. The shopping estimator converges in 542 CZs and the commuting estimator in 530; the comparison uses the 526 CZs with both. Because travel time enters in levels, I report the exact percentage decline per minute, $d^m_c = 100[1 - \exp(\widehat{\beta}^m_c)]$. Table A.14 reports the coefficient moments; the mean declines are 8.26% for shopping and 5.92% for commuting, and the cross-CZ correlation of the two coefficients is 0.86.

Table A.14: CZ-level shopping and commuting decay estimates

Flow	Converged CZ fits	Mean $\widehat{\beta}_c$	Median $\widehat{\beta}_c$	Median decline per minute
Shopping	542	−0.09	−0.08	8.08%
Commuting	530	−0.06	−0.06	6.10%

Notes: Coefficient moments and exact declines are computed on the 526-CZ matched comparison sample; the converged-fit column counts estimates available before matching the two modes.

The estimates are in line with the literature. Ahlfeldt et al. (2015) estimate a commuting semi-elasticity of about 0.07 per minute in Berlin, against the median 6.10% here, and Monte et al. (2018) estimate a county-to-county distance elasticity of −4.43. For consumption trips, Miyauchi et al. (2025) estimate travel-time elasticities of −1.82 from home and −1.81 from work in Tokyo, and Oh and Seo (2023) find that ten additional minutes reduce service-trip flows by about 80 percent in Seoul, against 57.00% here. The comparison is descriptive. Identification of β^m_c requires that travel time be uncorrelated with omitted pair-specific determinants of flows conditional on the fixed effects, and the visit data count trips rather than dollars. In the model the two coefficients map to different structural objects, $\beta^S_c = (1 - \sigma_N)\kappa$ and $\beta^C_c = -\chi\kappa$, which is why the comparison is made on the reduced-form declines rather than on the parameters.

A.6 Fact 3: Residential exposure designs

Fact 3 compares service employment in the tracts most and least exposed to entrant workers' spending. This section gives the alternative exposure measures and specifications estimated on the restricted data, the public analog built from predicted residences, and the reasons the observed

shopping network is not used as exposure.

The home-and-adjacent kernel. Every exposure measure in this section places each hire's spending at a home tract i and spreads it over $\mathcal{N}(i)$, the home tract and all 2020 tracts sharing a boundary point with it, with weights K_{ij} that assign 30 percent to the home tract and share 70 percent equally among the neighbors,

$$K_{ij} \geq 0, \qquad \sum_{j \in \mathcal{N}(i)} K_{ij} = 1, \qquad K_{ij} = 0 \text{ if } j \notin \mathcal{N}(i), \tag{A.10}$$

so each hire is counted once across tracts. The 30/70 split is the share of visits to the home tract among visits that end in the home tract or an adjacent tract (Appendix A.6). The entrant's workplace tract is excluded from every outcome sample.

Stacked new-income exposure. The first alternative ranks tracts by new income rather than by workers. New income is the earnings increase over the previous job for previously employed local hires and the full entrant wage for migrants and hires without a prior job. It is assigned to the hire's residence four quarters after hire, so that G_{ei} is new income received by hires living in tract i, and exposure is

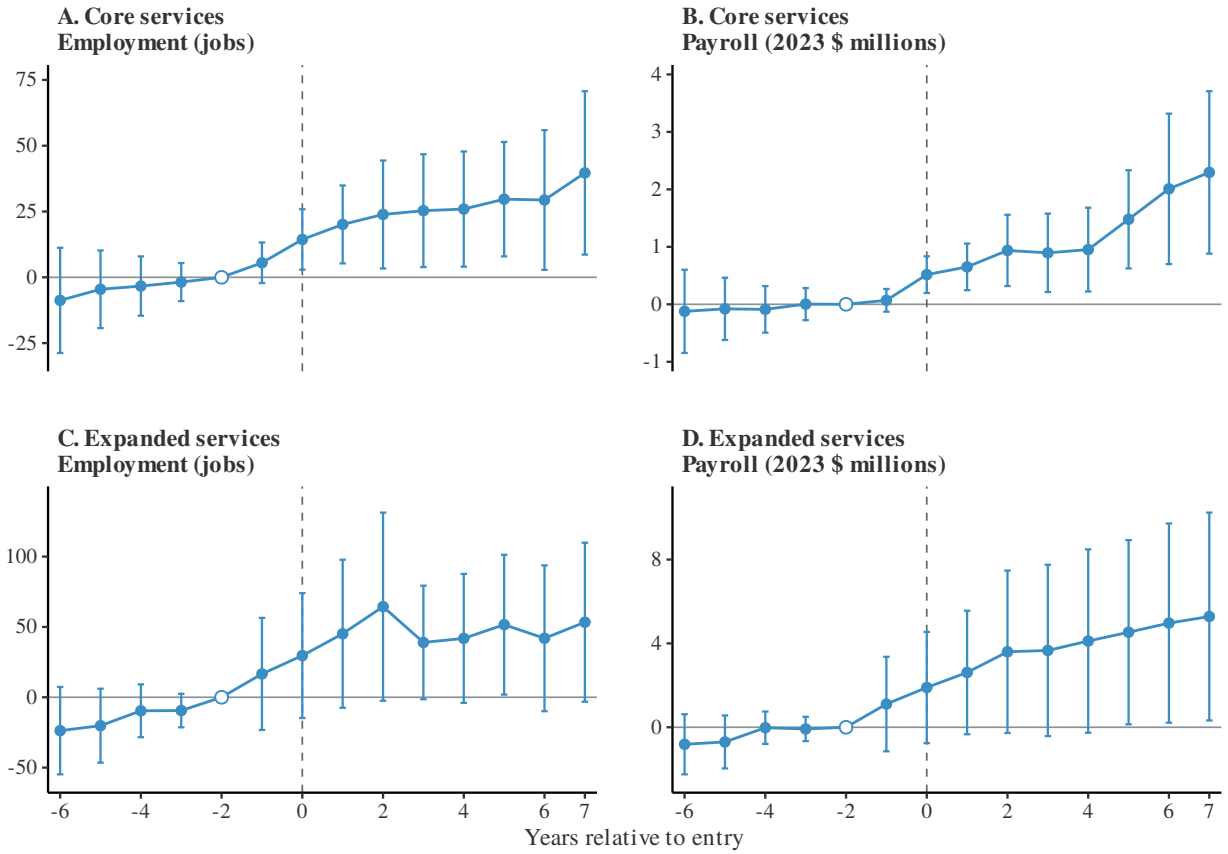
$$A^I_{ej} = \sum_i G_{ei} K_{ij}. \tag{A.11}$$

As in Fact 3, the sample is events in CZs with no other observed event, and I compare the top and bottom quartiles of A^I_{ej} within each event with equation (2). Mean exposure differs between the quartiles by \$1.166 million for the core-service outcomes (NAICS 44, 45, 722, 81) and \$1.164 million for the expanded outcomes that add 61 and 62; the release covers 30 CZs. By year seven, the core-service contrasts are 39.67 jobs and \$2.295 million in annual payroll, with both intervals above zero (Figure A.8). The expanded outcomes give 53.37 jobs and \$5.283 million, with the employment interval including zero. Pre-entry intervals include zero in all four panels. Because this measure uses post-hire residences, it can reflect moves after hiring as well as income gains at existing addresses.

Continuous exposure. The continuous specification uses all events and all tracts, including CZs with several events and the middle of the exposure distribution. For event e, let H_{ei} be the number of hires who lived in tract i of the event CZ before entry, with $A^W_{ej} = \sum_i H_{ei} K_{ij}$ the neighborhood allocation. Calendar-year exposure is $X^W_{jt} = 100^{-1} \sum_{e:\tau_e = t} A^W_{ej}$ for workers and $X^I_{jt} = 10^{-6} \sum_{e:\tau_e = t} A^I_{ej}$ for income. For either measure X and outcome o,

$$Y^o_{jt} = \alpha^o_j + \delta^o_{c(j)t} + \sum_{k \in \{-6,\ldots,7\} \setminus \{-2\}} \beta^o_k X_{j,t-k} + \gamma^o_- X^{\text{old}}_{jt} + \gamma^o_+ X^{\text{future}}_{jt} + \varepsilon^o_{jt}, \tag{A.12}$$

Figure A.8: Stacked service responses to neighborhood new-income exposure



Notes: Top-minus-bottom-quartile contrasts in new-income exposure from equation (2), in CZs with no other observed events. Panels A–B cover NAICS 44, 45, 722, and 81; panels C–D add 61 and 62. Income is assigned to residences four quarters after hire and spread with the 30/70 kernel. Bars are 95 percent confidence intervals clustered by CZ; year −2 is normalized to zero. Employment is in jobs and payroll in millions of 2023 dollars. 30 CZs.

with tract and CZ-by-year fixed effects, no weights, CZ-clustered standard errors, and stocks of exposure from events outside the window carried in X^{old} and X^{future}. Outcomes cover NAICS 44, 45, 722, and 81. Figure 5 in Section 3 reports the worker-exposure estimates. Stacked and continuous coefficients answer different questions. The stacked contrast is the change in the gap between the top and bottom quartiles in single-event CZs, while the continuous coefficient is a slope per 100 workers across the full distribution in all events, so their magnitudes are not directly comparable.

The public analog. The public design replaces identified hires with predicted residences and LEHD outcomes with LODES employment. It covers the 124 registry events dated 2011–2017 in 69 CZs.

Residences are predicted in two steps. First, for each event, the pre-entry LODES origin–destination counts $C_{hw_e,t}$ from residence tract h to the workplace tract w_e are averaged over event

years -4 through -2 and normalized within the CZ,

$$\overline{C}_{eh} = \frac{1}{3}\sum_{\ell=2}^{4} C_{hw_e,\tau_e-\ell}, \qquad p^C_{eh} = \frac{\overline{C}_{eh}}{\sum_{r\in c_e}\overline{C}_{er}}. \tag{A.13}$$

These are the residences of all workers in the tract, not of the entrant's hires. Second, because movers may settle outside the existing commuting catchment, a mover distribution p^M_{eh} is the share of the CZ's recent in-migrants aged 18–64 living in tract h, from the ACS five-year table B07001 ending in $\tau_e - 2$,

$$p^M_{eh} = \frac{M_{h,\tau_e-2}}{\sum_{r\in c_e} M_{r,\tau_e-2}}. \tag{A.14}$$

The predicted residence distribution mixes the two with a fixed weight,

$$p_{eh} = \omega\, p^C_{eh} + (1-\omega)\, p^M_{eh}, \qquad \omega = 0.84, \tag{A.15}$$

where ω is the share of CZ-resident hires who were previously employed or nonemployed in the CZ, and $1-\omega$ the share who migrated in, from the source shares behind Fact 1. The weights were fixed before any outcome was examined. Predicted residences are then scaled to the event's exposure mass N_e, the promised jobs allocated to CZ tracts by pre-entry commuting shares, on the support S_e that excludes the workplace tract and its neighbors,

$$R_{eh} = N_e\, \frac{p_{eh}\,\mathbf{1}\{h \in S_e\}}{\sum_{r\in S_e} p_{er}}. \tag{A.16}$$

Exposure then follows the home-and-adjacent rule. Let $\mathcal{A}_e(h)$ be the tracts in c_e adjacent to h. The 2019 visit data assign 0.09 of all service visits to the home tract and 0.20 to adjacent tracts; conditional on these local destinations, the home share is 0.31 and the adjacent share 0.69, which are the 30/70 weights used throughout. Destination exposure is

$$q^{HT}_{ehj} = 0.31\,\mathbf{1}\{j=h\} + \frac{0.69}{|\mathcal{A}_e(h)|}\,\mathbf{1}\{j\in\mathcal{A}_e(h)\}, \qquad E_{ej} = N_e\,\frac{\sum_h R_{eh}\, q^{HT}_{ehj}}{\sum_m\sum_h R_{eh}\, q^{HT}_{ehm}}, \tag{A.17}$$

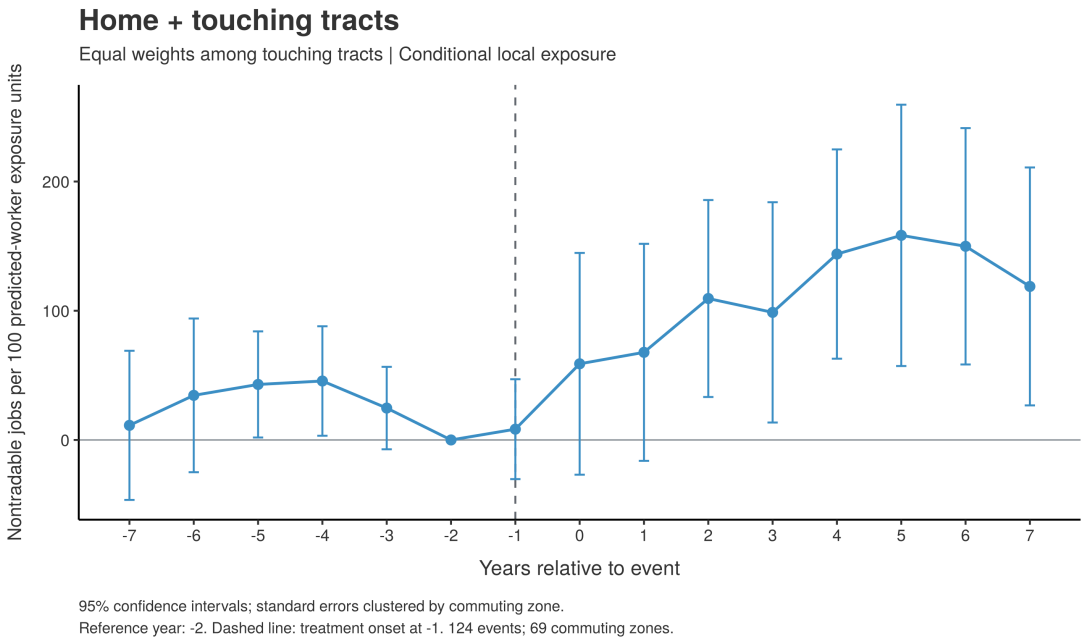
so each event's exposure mass is preserved. With $X^k_{jt} = \sum_e (E_{ej}/100)\,\mathbf{1}\{t-\tau_e = k\}$, the specification is

$$Y_{jt} = \alpha_j + \lambda_{c(j)t} + \sum_{k=-7,\,k\neq-2}^{7} \beta_k X^k_{jt} + \gamma_O O_{jt} + \gamma_F F_{jt} + \varepsilon_{jt}, \tag{A.18}$$

where Y_{jt} is LODES private service employment, α_j and $\lambda_{c(j)t}$ are tract and CZ-by-year fixed effects, and O_{jt} and F_{jt} carry exposure to events more than seven years in the past and future. The regression is unweighted and clustered by CZ, with 772,144 tract-years on 36,791 tracts over 2002–2022. The coefficient β_k is the differential employment response per 100 predicted-worker exposure units relative to year -2. Figure A.9 plots the estimates. The mean over years $+2$ through $+7$ is 129.90 jobs per 100 units (SE 41.00), and the joint test that the leads from -7 through -3 are zero

has $p = 0.06$, so the pre-entry evidence warrants caution. Like Fact 3, this is a within-CZ spatial response, not net job creation.

Figure A.9: Public home-plus-adjacent exposure and service employment



Notes: Coefficients β_k from equation (A.18), in private service jobs per 100 predicted-worker exposure units, with 95 percent confidence intervals clustered by CZ. Year -2 is normalized to zero; the dashed line marks treatment onset at -1. 124 events in 69 CZs.

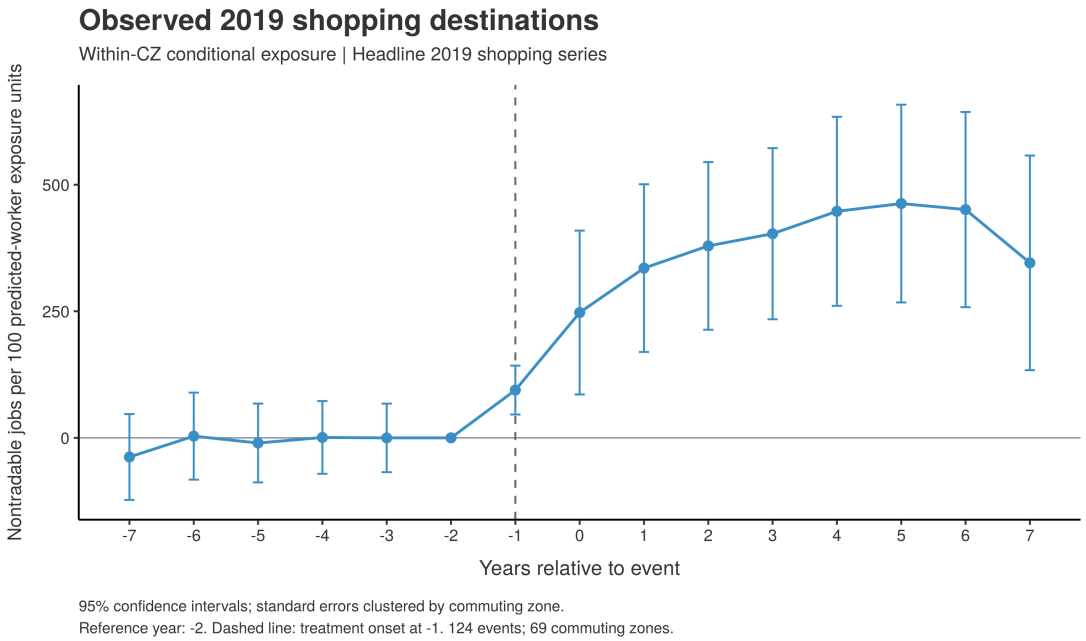
Why the observed shopping network is not used as exposure. The same predicted residences can instead be routed to destinations through the observed 2019 within-CZ visit shares, $S_{ej} = \sum_h R_{eh} s_{ehj}$, or through a gravity network built from pre-entry destination employment and travel time. Both alternatives have the same outcome, specification, and sample as the home-and-adjacent design. Two features rule them out as identifying variation. First, residents of different neighborhoods in a CZ shop at the same hubs, so exposure routed through the network lands on nearly the same destination tracts for every event in a CZ. Among the 85 entries in CZs with more than one entry, destination fixed effects alone explain 76.60% of the variation in network exposure. Second, the 2019 network is measured after the 2011–2017 events and may reflect their effects.

Figure A.10 shows what the observed network produces anyway. The mean coefficient over years +2 through +7 is 414.90 jobs per 100 predicted-worker exposure units (SE 89.20), several times the home-and-adjacent estimate, with a joint pretrend p-value of 0.37. The log and stacked versions of the same design reject the pretrend test ($p = 0.00$ and $p = 0.01$). Placebo networks that shuffle destination exposure within event, size-decile, and travel-time groups give a mean post-entry coefficient of 68.70 against the observed 414.90.[34] The coefficient measures employment growth at

[34]The placebo distribution uses 199 shuffles that preserve each event's exposure mass and destination support; a

the shared shopping destinations of each market in the years after entries, which is where entry-linked spending lands, but it cannot attribute that growth to a particular entrant. The residential design remains the identifying one.

Figure A.10: Observed 2019 shopping-destination exposure and service employment



Notes: Coefficients from equation (A.18) with exposure assigned through observed 2019 within-CZ shopping destination shares, in private service jobs per 100 predicted-worker exposure units, with 95 percent confidence intervals clustered by CZ. Year -2 is normalized to zero; the dashed line marks treatment onset at -1. 124 events in 69 CZs. Because the 2019 network postdates the events, the series is descriptive.

A.7 Aggregate and median housing shares across commuting zones

The aggregate housing share measures the fraction of a CZ's total household resources devoted to housing. The median household share describes the housing expenditure burden of the middle household. Figure A.11 compares these moments using a common sample and measurement procedure across 583 CZs in the 50 states and the District of Columbia. This national comparison uses a common 2008–2017 period; the event calibration in equation (C.14) targets the household median in each CZ's pre-entry sample.

Sample and geographic assignment. The data are the Census Bureau's national ACS five-year Public Use Microdata Samples for 2008–2012 and 2013–2017, linking housing records to their adult person records. Each release receives one half of its published survey weight. Pooling these nonoverlapping releases produces a ten-year household distribution and average population stocks. The housing sample contains 12,113,453 occupied households with at least one adult aged 18 or

restricted wild-cluster bootstrap with 999 CZ-level draws gives $p = 0.00$.

older and usable housing-service measures. Rents and incomes use the ACS adjustment factors followed by the CPI conversion to annual 2023 dollars.

Public Use Microdata Areas (PUMAs) are allocated fractionally to CZs using 2010 census population in their intersections. The allocation uses 2000 PUMAs for 2008–2011 interviews and 2010 PUMAs for 2012–2017 interviews. The former crosswalk is constructed from the Missouri Census Data Center's Geocorr PUMA–county intersections; the latter uses Census tract–PUMA relationships and tract populations. Historical county codes and merged PUMAs are reconciled before aggregation. For household m in PUMA p, let a_{pc} be the fraction allocated to CZ c. Its estimation weight is $w_{mc} = \frac{1}{2}\text{WGTP}_m a_{pc}$.

Housing services and household resources. Paying renters contribute their annual contract rent. Owner-occupied and rent-free dwellings contribute imputed rental services from a national, survey-weighted regression estimated on 3,322,577 paying renter households:

$$\log E_m^{\text{rent}} = X_m'\beta + \lambda_{p(m),t(m)} + u_m. \tag{A.19}$$

Here X_m contains log rooms and its square, bedrooms and its square, structure-type indicators, and structure-type interactions with log rooms. The fixed effects identify state, PUMA vintage, PUMA, and interview year. Imputed annual service expenditure is $E_m = \exp(X_m'\widehat{\beta} + \widehat{\lambda}_{p(m),t(m)})\widehat{s}_{b(m)}$, where $\widehat{s}_b$ is the survey-weighted mean of $\exp(\widehat{u}_m)$ among renters of structure type b. This retransformation incorporates the residual distribution when converting predicted log rents to levels.

Household resources sum income and housing-portfolio dividends over adults. An employed adult receives 95 percent of observed wage and salary income. A nonworker receives the CZ–skill mean of public assistance, Supplemental Security Income, and Social Security income among nonworkers, estimated with person weights. High skill denotes a bachelor's degree or higher. The CZ housing portfolio equals total housing-service expenditure divided by 1.28. Its skill shares are based on reported owner-occupied housing values, split equally between the adult householder and spouse or unmarried partner when both are present. Each skill's share of the portfolio is distributed equally among its represented adults in the eligible sample. Thus

$$Y_{mc} = Y_{mc}^{\text{inc}} + \sum_{k\in\{H,L\}} n_{mk} D_{ck}, \qquad D_{ck} = \frac{\omega_{ck}}{N_{ck}} \frac{\sum_m w_{mc} E_m}{1.28}, \tag{A.20}$$

where n_{mk} is the household's adult count of skill k, $N_{ck} = \sum_m w_{mc} n_{mk}$, and ω_{ck} is that skill's ownership-value share. Property tax revenue leaves the modeled economy, so residents receive no rebate. These resources implement the model's individual budget convention.

The two expenditure shares. Both statistics use the same eligible households, service bills, and resources. With $s_{mc} = E_m/Y_{mc}$, define

$$A_c = \frac{\sum_m w_{mc} E_m}{\sum_m w_{mc} Y_{mc}} = \sum_m \frac{w_{mc} Y_{mc}}{\sum_a w_{ac} Y_{ac}} s_{mc}, \tag{A.21}$$

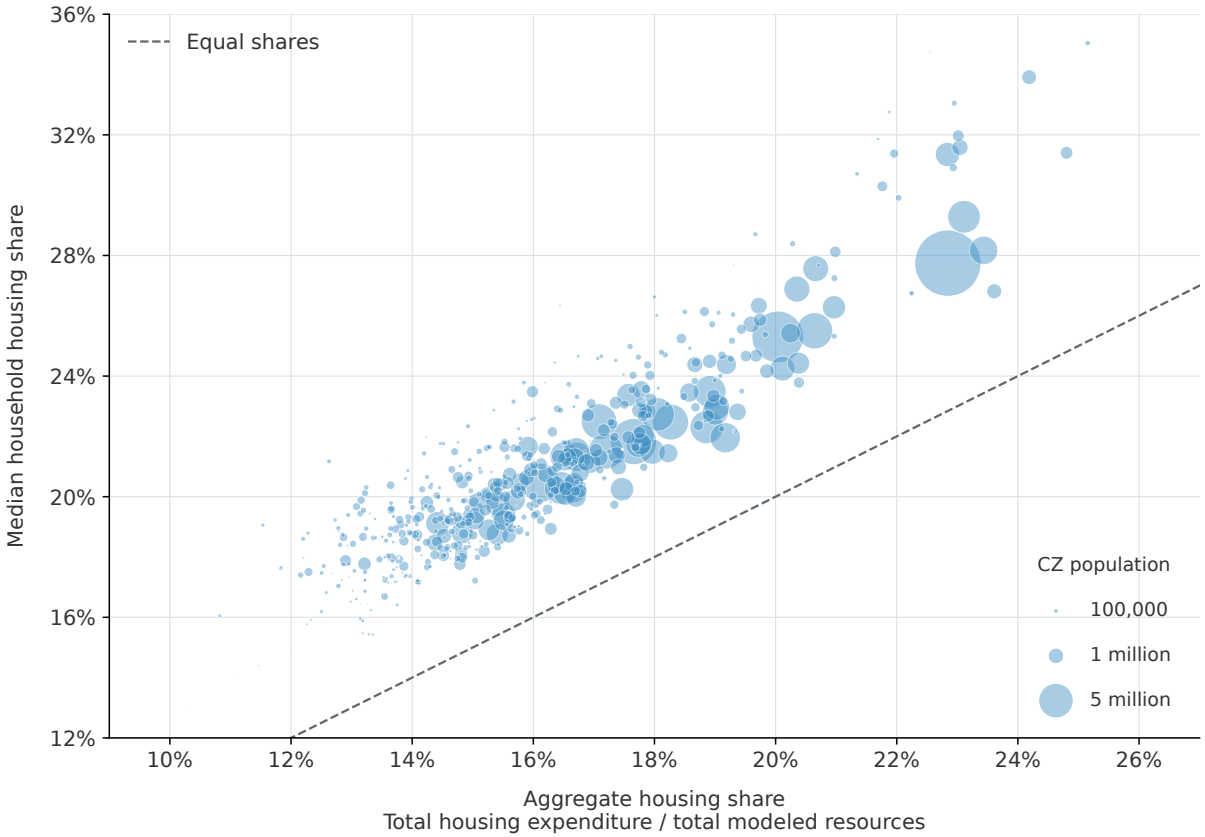
$$M_c = \inf \left\{ s : \frac{\sum_m w_{mc} \mathbf{1}\{s_{mc} \leq s\}}{\sum_m w_{mc}} \geq \frac{1}{2} \right\}. \tag{A.22}$$

The aggregate A_c weights household shares by resources as well as survey weights; M_c is the household-weighted median of the pooled observations. The aggregate share describes housing's claim on total resources, while the median describes a typical household's expenditure burden. Greater resource weight on households with lower housing shares can therefore produce a lower aggregate share.

Across CZs, the median of A_c is 15.40 percent and the median of M_c is 20.20 percent. The median within-CZ gap is 4.70 percentage points, and the household median exceeds the aggregate in all 583 CZs. Each dot's area is proportional to average 2008–2017 population, obtained from all-age person records, including group quarters, with the same half-weight pooling and geographic allocation. The cross-CZ summaries give each CZ equal weight.

The comparison isolates the choice of expenditure-share statistic under common measurement. Both moments depend on the rental-service imputation and resource convention. The model targets the median household share; its baseline portfolio and housing quantities are determined jointly by the budget and market-clearing equations. A common Cobb–Douglas share generally cannot reproduce both the observed median and aggregate share. Fractional geographic assignment assumes that households in a PUMA represent its constituent CZ portions; small CZs sharing a PUMA can consequently have similar estimates. The plotted values are point estimates conditional on these imputations and geographic allocations.

Figure A.11: Aggregate and median household housing shares across U.S. commuting zones



Notes: ACS 2008–2012 and 2013–2017 five-year microdata, pooled with one half of each release's survey weights. The horizontal axis is total housing-service expenditure divided by total modeled household resources; the vertical axis is the household-weighted median expenditure share. Both statistics use identical eligible households and measurement rules. Housing services are observed contract rents or imputed rental equivalents. Resources include modeled income and housing-portfolio dividends. Dot area represents average all-age CZ population during 2008–2017. The dashed line denotes equal shares. All 583 CZs are shown; the axes have different ranges. PUMA-to-CZ locations are assigned using population weights.

B Model

This appendix derives individual demand, employment and location choices, firms' input demands, and the equilibrium accounts from the preferences and technologies in Section 4. A superscript 0 denotes the no-entry equilibrium. Appendix C describes the quantification and solution of these conditions.

B.1 Individual preferences

An individual ω of skill k chooses consumption and a joint residence and employment choice ℓ. For workers, $\ell = (i, j, g)$ specifies residence, workplace, and sector. For nonworkers it specifies residence and sector. The housing share is $\alpha_\ell = \alpha_H$ for $i \in \mathcal{J}$ and $\alpha_\ell = \overline{\alpha}_H^o$ for $i = o$. Write $\gamma_\ell = (1 - \alpha_\ell)s_N$

and $\beta_\ell = (1-\alpha_\ell)(1-s_N)$, so that $\alpha_\ell + \gamma_\ell + \beta_\ell = 1$. The utility in (3) can be written

$$U_\ell(\omega) = z_\ell(\omega)\mathcal{B}_\ell h^{\alpha_\ell} n^{\gamma_\ell} x^{\beta_\ell}, \qquad \mathcal{B}_\ell = \begin{cases} B^R_{ik} B^W_{jgk}/d_{ij}, & i,j \in \mathcal{J},\ e = W, \\ \overline{B}^{\text{out}}_{igk}, & i \in \mathcal{J},\ j = o,\ e = W, \\ \overline{B}^{\text{in}}_{jgk}, & i = o,\ j \in \mathcal{J},\ e = W, \\ \overline{B}^{\text{oo}}_{gk}, & i = j = o,\ e = W, \\ B^R_{ik} B^U_{ik}, & e = U. \end{cases} \tag{B.1}$$

Here $e = W$ denotes work and $e = U$ nonemployment. For internal commuters, $d_{ij} = \exp(\kappa t_{ij})$. The boundary amenity composites include fixed commuting costs and the attractiveness of access to outside jobs or residences. Each has at most one location index within the CZ; they preserve the separately calibrated inbound and outbound flows and the outside worker stock. Their construction is given in Appendix C.6.1. Internal commuting retains separate residence and workplace factors, with workplace amenities varying by sector. Housing and the two nonhousing goods have constant expenditure shares. After optimizing consumption and absorbing fixed expenditure-share constants into participation preference parameters, the systematic value is $\Phi_\ell = \mathcal{B}_\ell y_\ell / [(q^R_i)^{\alpha_\ell} \mathcal{P}^{\gamma_\ell}_\ell]$. The next subsection derives the employment and location probabilities; Appendix B.3 derives income, consumption demands, and this normalization.

B.2 Nested tastes and employment and location choices

Individuals choose jointly among residence and employment choices. The nested Fréchet distribution allows their tastes to be more similar among location pairs than across participation choices, sectors, and areas. Draws are independent across individuals and are held fixed across equilibria. For a nest v containing choices $a \in \mathcal{C}(v)$, let $\overline{B}_{va,k}$ be preference parameters and θ_v its shape parameter. The value index and conditional choice probability are

$$V_{vk} = \left[\sum_{a \in \mathcal{C}(v)} \overline{B}_{va,k} V^{\theta_v}_{ak} \right]^{1/\theta_v}, \tag{B.2}$$

$$\pi_{a|v,k} = \frac{\overline{B}_{va,k} V^{\theta_v}_{ak}}{V^{\theta_v}_{vk}}. \tag{B.3}$$

For a specific residence and employment choice, $V_{ak} = \Phi_\ell$ and $\overline{B}_{va,k} = 1$: fixed location attractiveness is contained in the amenities in Φ_ℓ. The $\overline{B}$'s describe preferences over broader choices, whose values summarize the residence and employment opportunities they include. A nest with one potential choice has a preference parameter equal to one. The order of the region and sector nests differs by skill. Low-skill individuals choose their region before their sector, and high-skill individuals choose their sector before their region, so the shape parameters satisfy

$$0 < \mu_L \leq \nu \leq \eta \leq \psi \leq \chi, \qquad 0 < \nu \leq \mu_H \leq \eta \leq \psi \leq \chi. \tag{B.4}$$

The migration parameter μ_k differs by skill, while the sector and participation parameters ν and η are common across skills and CZs. In the calibration, $\mu_L = 1.51$, $\nu = 1.64$, $\mu_H = 2.32$, and $\eta = 3.32$, so both orderings hold. The residence–workplace parameter χ is estimated from commuting gravity for each CZ, and ψ varies across events, so the ordering holds event by event. Their calibration is reported in Appendix C.1.

B.2.1 Value indices and conditional shares

The region consists of two areas, the CZ and the accessible outside area. Let $r \in \{I, O\}$ denote the area of residence, so that $r = I$ for CZ residents and $r = O$ for outside-area residents. Let $g \in \mathcal{G} = \mathcal{G}_T \cup \mathcal{G}_N$ be a sector, with $\mathcal{G}_T = \{T_H, T_L\}$ and $\mathcal{G}_N = \{N_H, N_L\}$ (for a nonworker, the sector in which they would work), and $e \in \{W, U\}$ work or nonemployment. A location is a CZ tract or the outside area, and a location pair ℓ is a residence and workplace location, or a residence location for a nonworker. CZ residents choose directly among location pairs with workplaces inside and outside the CZ. Outside-area residents who work first choose a job location $s \in \{I, O\}$, inside or outside the CZ, and then a location pair within it. Each nest's index equals the value index in (B.2) raised to the nest's shape parameter. The working, nonworking, sector, and area indices are defined in (B.8) and (B.9) below, and applying (B.3) at each margin gives the conditional shares that follow them. For low-skill individuals, residence in the region and in the rest of the country form the top nest,

$$(V_L^{\text{all}})^{\mu_L} = \overline{B}_L^{\text{nat}} (V_L^{\text{nat}})^{\mu_L} + \overline{B}_L^{\text{reg}} (V_L^{\text{reg}})^{\mu_L}, \qquad (V_L^{\text{reg}})^{\mu_L} = \sum_{r \in \{I,O\}} \overline{B}_{rL}^{R} \mathcal{Q}_{rL}^{\mu_L/\nu}, \tag{B.5}$$

where the national value V_L^{nat} is fixed for national population $\mathcal{N}_L$. Because migration and the choice between the two areas share μ_L, the choice to live in the region and the choice of area combine into one upper nest. High-skill individuals instead choose a sector at the top and then a region within it; the paragraph on the high-skill ordering in Appendix B.2.2 gives their indices and allocations.

For a group of location pairs $\mathcal{D}$ with shape χ,

$$\pi_{\ell|\mathcal{D},k} = \frac{\Phi_\ell^\chi}{\sum_{a \in \mathcal{D}} \Phi_a^\chi}. \tag{B.6}$$

For outside-area workers, multiplying this location-pair share by the job-location share yields (B.21). CZ residents and nonworkers have no job-location choice, and nonworkers choose among residence locations with shape χ.

B.2.2 Expanded utilities and population allocations

This subsection writes the indices and allocations explicitly for both areas and then derives the CZ resident notation used in the main text. The preference parameters are $\overline{B}_{rk}^R$ for area r, $\overline{B}_{rgk}^S$ for sector g, $\overline{B}_{rgk}^E$ and $\overline{B}_{rgk}^U$ for work and nonemployment, and $\overline{B}_{gsk}^O$ for job location s, a choice only

outside-area residents make. Write $\mathcal{L}^W_{Igk}$ for the working location pairs available to CZ residents, $\mathcal{L}^W_{Ogsk}$ for those available to outside-area residents with job location s, with $\mathcal{L}^W_{Ogk} = \mathcal{L}^W_{OgIk} \cup \mathcal{L}^W_{OgOk}$, and $\mathcal{L}^U_{rgk}$ for the nonworking residence locations. Let i_ℓ denote the residence location in pair ℓ and, for workers, j_ℓ its workplace location. The residence location determines the area, with $i_\ell \in \mathcal{I}_I = \mathcal{J}$ for CZ residents and $i_\ell \in \mathcal{I}_O = \{o\}$ for outside-area residents. Nonworking choices retain sector g, but their systematic utility and conditional residence probabilities are identical across sectors.

Working and nonworking utility. Using the housing share α_ℓ defined in Appendix B.1, systematic utility is

$$\Phi_\ell = \begin{cases} \dfrac{\mathcal{B}_\ell y_\ell}{(q^R_{i_\ell})^{\alpha_\ell}(P^{\rho}_{i_\ell} P^{1-\rho}_{j_\ell})^{(1-\alpha_\ell)s_N}}, & \text{if working,} \\ \dfrac{\mathcal{B}_\ell y_\ell}{(q^R_{i_\ell})^{\alpha_\ell} P^{(1-\alpha_\ell)s_N}_{i_\ell}}, & \text{if nonemployed.} \end{cases} \tag{B.7}$$

Disposable income y_ℓ is given by (B.37). Outside income and prices are fixed except that an outside resident working inside the CZ receives the endogenous wage at that workplace. The values use the expenditure-share normalization derived in Appendix B.3.

Working, nonworking, and sector indices. The index $\mathcal{W}_{Igk}$ summarizes CZ residents' working location pairs and $\mathcal{W}_{Ogsk}$ those of outside-area residents with job location s. The index $\widetilde{\mathcal{W}}_{Ogk}$ combines the two job locations, and $\mathcal{U}_{rgk}$ summarizes nonworking residence locations,

$$\begin{aligned} \mathcal{W}_{Igk} &= \sum_{\ell \in \mathcal{L}^W_{Igk}} \Phi^{\chi}_\ell, & \mathcal{W}_{Ogsk} &= \sum_{\ell \in \mathcal{L}^W_{Ogsk}} \Phi^{\chi}_\ell, \\ \widetilde{\mathcal{W}}_{Ogk} &= \sum_{s \in \{I,O\}} \overline{B}^O_{gsk} \mathcal{W}^{\psi/\chi}_{Ogsk}, & \mathcal{U}_{rgk} &= \sum_{\ell \in \mathcal{L}^U_{rgk}} \Phi^{\chi}_\ell, \end{aligned} \tag{B.8}$$

with Φ_ℓ from (B.7). Working and nonworking opportunities combine within sector, and sectors combine within each area,

$$\begin{aligned} \mathcal{S}_{Igk} &= \overline{B}^E_{Igk} \mathcal{W}^{\eta/\chi}_{Igk} + \overline{B}^U_{Igk} \mathcal{U}^{\eta/\chi}_{Igk}, \\ \mathcal{S}_{Ogk} &= \overline{B}^E_{Ogk} \widetilde{\mathcal{W}}^{\eta/\psi}_{Ogk} + \overline{B}^U_{Ogk} \mathcal{U}^{\eta/\chi}_{Ogk}, \\ \mathcal{Q}_{rk} &= \sum_{g \in \mathcal{G}} \overline{B}^S_{rgk} \mathcal{S}^{\nu/\eta}_{rgk}. \end{aligned} \tag{B.9}$$

Because CZ residents have no job-location choice, ψ does not enter their indices.

CZ resident notation. Dropping the area subscript for CZ residents gives the main-text notation,

$$\begin{aligned} \mathcal{W}_{gk} &= \mathcal{W}_{Igk}, & \mathcal{U}_{gk} &= \mathcal{U}_{Igk}, & \mathcal{S}_{gk} &= \mathcal{S}_{Igk}, \\ \overline{B}^{E}_{gk} &= \overline{B}^{E}_{Igk}, & \overline{B}^{U}_{gk} &= \overline{B}^{U}_{Igk}, & \overline{B}^{S}_{gk} &= \overline{B}^{S}_{Igk}, \\ N_k &= N_{Ik}, & N_{gk} &= N_{Igk}, & E_{gk} &= E_{Igk}, \end{aligned} \tag{B.10}$$

with choice sets $\mathcal{L}^{W}_{gk} = \mathcal{L}^{W}_{Igk}$ and $\mathcal{L}^{U}_{gk} = \mathcal{L}^{U}_{Igk}$. The first line of (B.9) is then the main-text sector index (7).

Regional migration and residence. This paragraph and the next describe the low-skill ordering, so $k = L$ and the migration parameter is μ_L. The fixed national contribution is $\overline{\mathcal{O}}_k = \overline{B}^{\text{nat}}_k (V^{\text{nat}}_k)^{\mu_L}$, and V^{reg}_k is given in (B.5). The regional share and the share living in each area are

$$\begin{aligned} b^{\text{reg}}_k &= \frac{\overline{B}^{\text{reg}}_k \sum_{r\in\{I,O\}} \overline{B}^{R}_{rk} \mathcal{Q}^{\mu_L/\nu}_{rk}}{\overline{\mathcal{O}}_k + \overline{B}^{\text{reg}}_k \sum_{r\in\{I,O\}} \overline{B}^{R}_{rk} \mathcal{Q}^{\mu_L/\nu}_{rk}}, \\ \pi(r \mid \text{reg}, k) &= \frac{\overline{B}^{R}_{rk} \mathcal{Q}^{\mu_L/\nu}_{rk}}{\sum_{r'\in\{I,O\}} \overline{B}^{R}_{r'k} \mathcal{Q}^{\mu_L/\nu}_{r'k}}. \end{aligned} \tag{B.11}$$

Multiplying these shares gives

$$\begin{aligned} N_{rk} &= \mathcal{N}_k b^{\text{reg}}_k \pi(r \mid \text{reg}, k) = \mathcal{N}_k \frac{\overline{B}^{\text{reg}}_k \overline{B}^{R}_{rk} \mathcal{Q}^{\mu_L/\nu}_{rk}}{\overline{\mathcal{O}}_k + \overline{B}^{\text{reg}}_k \sum_{r'\in\{I,O\}} \overline{B}^{R}_{r'k} \mathcal{Q}^{\mu_L/\nu}_{r'k}}, \\ N^{\text{reg}}_k &= \sum_r N_{rk} = \mathcal{N}_k b^{\text{reg}}_k, \qquad N^{\text{nat}}_k = \mathcal{N}_k (1 - b^{\text{reg}}_k). \end{aligned} \tag{B.12}$$

Both migration and the choice between the two areas use μ_L. For the main-text equation, define

$$\begin{aligned} \overline{B}^{R}_{L} &= \overline{B}^{\text{reg}}_{L} \overline{B}^{R}_{IL}, \\ \mathcal{O}_L &= \overline{\mathcal{O}}_L + \overline{B}^{\text{reg}}_{L} \overline{B}^{R}_{OL} \mathcal{Q}^{\mu_L/\nu}_{OL}. \end{aligned} \tag{B.13}$$

The term $\mathcal{O}_L$ includes both the fixed national contribution and living in the outside area. It is endogenous because outside-area residents can work at CZ workplaces. Setting $r = I$ in (B.12) and substituting $\mathcal{Q}_{IL} = \mathcal{Q}_L = \sum_g \overline{B}^{S}_{gL} \mathcal{S}^{\nu/\eta}_{gL}$ gives the CZ's low-skill population N_L.

Sector allocation. The share of area-r residents choosing sector g is $\overline{B}^{S}_{rgk} \mathcal{S}^{\nu/\eta}_{rgk} / \mathcal{Q}_{rk}$, so

$$N_{rgk} = N_{rk} \frac{\overline{B}^{S}_{rgk} \mathcal{S}^{\nu/\eta}_{rgk}}{\mathcal{Q}_{rk}}. \tag{B.14}$$

For $r = I$, this gives the low-skill case of (8). The allocation includes both employed and nonemployed individuals. Nonemployment income, prices, and amenities are identical across sectors at each residence location. The baseline allocation and its normalization are derived below.

High-skill ordering. High-skill individuals choose a sector first and then a region within it. Let $\overline{B}^S_{gH}$ be the preference parameter for sector g, common to both areas, $\overline{B}^{\text{nat}}_{gH}$ and $\overline{B}^{\text{reg}}_{gH}$ those for living elsewhere in the country or in the region within sector g, and $\overline{B}^R_{rgH}$ that for area r. The value of sector g combines the fixed national value of that sector elsewhere in the country, V^{nat}_{gH}, with the two areas,

$$\begin{aligned} \mathcal{V}_{gH} &= \overline{\mathcal{O}}_{gH} + \overline{B}^{\text{reg}}_{gH} \sum_{r\in\{I,O\}} \overline{B}^R_{rgH} \mathcal{S}^{\mu_H/\eta}_{rgH}, \qquad & \overline{\mathcal{O}}_{gH} &= \overline{B}^{\text{nat}}_{gH} (V^{\text{nat}}_{gH})^{\mu_H}, \\ (V^{\text{all}}_H)^{\nu} &= \sum_{g\in\mathcal{G}} \overline{B}^S_{gH} \mathcal{V}^{\nu/\mu_H}_{gH}, \end{aligned} \tag{B.15}$$

with $\mathcal{S}_{rgH}$ from (B.9). The sector share, the regional share within the sector, and the area share are

$$\begin{aligned} \pi(g \mid H) &= \frac{\overline{B}^S_{gH} \mathcal{V}^{\nu/\mu_H}_{gH}}{\sum_{g'\in\mathcal{G}} \overline{B}^S_{g'H} \mathcal{V}^{\nu/\mu_H}_{g'H}}, \\ b^{\text{reg}}_{gH} &= \frac{\overline{B}^{\text{reg}}_{gH} \sum_r \overline{B}^R_{rgH} \mathcal{S}^{\mu_H/\eta}_{rgH}}{\mathcal{V}_{gH}}, \qquad & \pi(r \mid \text{reg}, g, H) &= \frac{\overline{B}^R_{rgH} \mathcal{S}^{\mu_H/\eta}_{rgH}}{\sum_{r'} \overline{B}^R_{r'gH} \mathcal{S}^{\mu_H/\eta}_{r'gH}}, \end{aligned} \tag{B.16}$$

so the numbers by area and sector are

$$\begin{aligned} N_{rgH} &= \mathcal{N}_H \, \pi(g \mid H) \, b^{\text{reg}}_{gH} \, \pi(r \mid \text{reg}, g, H) = \mathcal{N}_H \, \pi(g \mid H) \, \frac{\overline{B}^{\text{reg}}_{gH} \overline{B}^R_{rgH} \mathcal{S}^{\mu_H/\eta}_{rgH}}{\mathcal{V}_{gH}}, \\ N_{rH} &= \sum_{g\in\mathcal{G}} N_{rgH}, \qquad N^{\text{nat}}_H = \mathcal{N}_H \sum_{g\in\mathcal{G}} \pi(g \mid H)\big(1 - b^{\text{reg}}_{gH}\big). \end{aligned} \tag{B.17}$$

For the main-text equation, define

$$\overline{B}^R_{gH} = \overline{B}^{\text{reg}}_{gH} \overline{B}^R_{IgH}, \qquad \mathcal{O}_{gH} = \overline{\mathcal{O}}_{gH} + \overline{B}^{\text{reg}}_{gH} \overline{B}^R_{OgH} \mathcal{S}^{\mu_H/\eta}_{OgH}, \tag{B.18}$$

so that $\mathcal{V}_{gH} = \mathcal{O}_{gH} + \overline{B}^R_{gH} \mathcal{S}^{\mu_H/\eta}_{gH}$. Setting $r = I$ in (B.17) gives the high-skill case of (8). As for low-skill individuals, $\mathcal{O}_{gH}$ is endogenous because outside-area residents can work at CZ workplaces. Each sector has its own national option, and within sector g the migration elasticity is $\mu_H(1 - b^{\text{reg}}_{gH})$.

Labor-force participation. The employed share of those choosing sector g is the working term of $\mathcal{S}_{rgk}$ relative to its total,

$$\frac{E_{Igk}}{N_{Igk}} = \frac{\overline{B}^{E}_{Igk}\mathcal{W}^{\eta/\chi}_{Igk}}{\mathcal{S}_{Igk}}, \qquad \frac{E_{Ogk}}{N_{Ogk}} = \frac{\overline{B}^{E}_{Ogk}\widetilde{\mathcal{W}}^{\eta/\psi}_{Ogk}}{\mathcal{S}_{Ogk}}. \tag{B.19}$$

The first gives (10). The nonemployed population in sector g is $N_{rgk} - E_{rgk}$.

Job location and location pairs. An outside-area resident who works in sector g chooses job location s and then a location pair $\ell \in \mathcal{L}^{W}_{Ogsk}$, with shares

$$\pi(s \mid O, g, W, k) = \frac{\overline{B}^{O}_{gsk}\mathcal{W}^{\psi/\chi}_{Ogsk}}{\widetilde{\mathcal{W}}_{Ogk}}, \qquad \pi_{\ell|O,g,W,s,k} = \frac{\Phi^{\chi}_{\ell}}{\mathcal{W}_{Ogsk}}. \tag{B.20}$$

A CZ resident chooses a location pair directly. The conditional location-pair probabilities of workers are therefore

$$\pi_{\ell|I,g,W,k} = \frac{\Phi^{\chi}_{\ell}}{\mathcal{W}_{Igk}}, \qquad \pi_{\ell|O,g,W,k} = \frac{\overline{B}^{O}_{gsk}\mathcal{W}^{\psi/\chi}_{Ogsk}}{\widetilde{\mathcal{W}}_{Ogk}}\frac{\Phi^{\chi}_{\ell}}{\mathcal{W}_{Ogsk}}, \tag{B.21}$$

and the main-text probability abbreviates $\pi_{\ell k} = \pi_{\ell|I,g,W,k}$. Multiplying the area–sector population N_{rgk}, from (B.14) for low-skill and (B.17) for high-skill individuals, by the employment and location shares gives the working population, for CZ residents and outside-area residents respectively,

$$\begin{aligned} F_{ijgk} &= N_{Igk}\underbrace{\frac{\overline{B}^{E}_{Igk}\mathcal{W}^{\eta/\chi}_{Igk}}{\mathcal{S}_{Igk}}}_{\text{work}}\underbrace{\frac{\Phi^{\chi}_{\ell}}{\mathcal{W}_{Igk}}}_{\text{location pair}}, \qquad i \in \mathcal{J}, \\ F_{ojgk} &= N_{Ogk}\underbrace{\frac{\overline{B}^{E}_{Ogk}\widetilde{\mathcal{W}}^{\eta/\psi}_{Ogk}}{\mathcal{S}_{Ogk}}}_{\text{work}} \\ &\quad \times \underbrace{\frac{\overline{B}^{O}_{gsk}\mathcal{W}^{\psi/\chi}_{Ogsk}}{\widetilde{\mathcal{W}}_{Ogk}}}_{\text{job location}}\underbrace{\frac{\Phi^{\chi}_{\ell}}{\mathcal{W}_{Ogsk}}}_{\text{location pair}}. \end{aligned} \tag{B.22}$$

The first line gives (11). Substituting the area–sector populations gives the unconditional number choosing each pair.

Nonemployment and population accounting. For a nonworking choice $\ell \in \mathcal{L}^U_{rgk}$ with residence location i and sector g, the nonemployed population is

$$U_{igk} = (N_{rgk} - E_{rgk})\frac{\Phi^\chi_\ell}{\mathcal{U}_{rgk}}, \qquad U_{ik} = \sum_{g\in\mathcal{G}} U_{igk}. \tag{B.23}$$

An unavailable nonworking choice has $U_{igk} = 0$. Thus $\mathcal{U}_{rgk}$ is a utility index, while U_{igk} and U_{ik} count individuals. The allocations satisfy

$$\begin{aligned} \sum_g N_{rgk} &= N_{rk}, & \sum_{\ell\in\mathcal{L}^W_{rgk}} F_{i_\ell j_\ell gk} &= E_{rgk}, \\ \sum_{i\in\mathcal{I}_r} U_{ik} &= \sum_g (N_{rgk} - E_{rgk}), & \sum_r N_{rk} + N^{\text{nat}}_k &= \mathcal{N}_k. \end{aligned} \tag{B.24}$$

Conditional shares apply to nests with positive value; an empty nest has zero population. Write $\pi^{\text{all}}_{\ell k}$ for the unconditional probability, distinguished from the conditional location-pair probability $\pi_{\ell k}$ in the main text. For a regional choice,

$$\begin{aligned} \pi^{\text{all}}_{\ell L} &= b^{\text{reg}}_L \pi(r \mid \text{reg}, L)\pi(g \mid r, L)\pi(e \mid r, g, L)\pi_{\ell|r,g,e,L}, \\ \pi^{\text{all}}_{\ell H} &= \pi(g \mid H)\, b^{\text{reg}}_{gH} \pi(r \mid \text{reg}, g, H)\pi(e \mid r, g, H)\pi_{\ell|r,g,e,H}. \end{aligned}$$

For workers the last factor is (B.21); for nonworkers it is the residence-location share in (B.23). The national choice has probability $1 - b^{\text{reg}}_L$ for low-skill and $\sum_g \pi(g \mid H)(1 - b^{\text{reg}}_{gH})$ for high-skill individuals. For the full feasible choice set $\mathcal{A}_k$,

$$n_{\ell k} = \mathcal{N}_k \pi^{\text{all}}_{\ell k}, \qquad \sum_{\ell\in\mathcal{A}_k} \pi^{\text{all}}_{\ell k} = 1. \tag{B.25}$$

These probabilities determine equilibrium populations. Individual transitions across equilibria are derived in Appendix B.12.

B.2.3 Derivation from the taste distribution

For any positive vector of choice values x, let $\mathcal{V}_k(x)$ denote the all-location value obtained from the indices above, with each location-pair value Φ_ℓ replaced by x_ℓ and the national value replaced by x_{nat}. Define $G_k(x) = \mathcal{V}_k(x)^{\theta^{\text{all}}_k}$, where the shape of the top nest is $\theta^{\text{all}}_L = \mu_L$ and $\theta^{\text{all}}_H = \nu$. The nested Fréchet joint distribution in (5) is therefore

$$\Pr\{z_\ell \le t_\ell \text{ for every } \ell\} = \exp\left\{ - \left[\mathcal{V}_k\big((t^{-1}_\ell)_\ell\big)\right]^{\theta^{\text{all}}_k} \right\}. \tag{B.26}$$

The ordering in (B.4) ensures that this is a valid nested extreme-value distribution. Since each value index is homogeneous of degree one, maximized utility has distribution

$$\Pr\{\max_{\ell} z_\ell \Phi_\ell \le u\} = \exp\left[-\left(\frac{V_k^{\text{all}}}{u}\right)^{\theta_k^{\text{all}}}\right]. \tag{B.27}$$

Let $\partial_\ell G_k$ denote the derivative of G_k with respect to x_ℓ. The probability that ℓ maximizes utility integrates the density that its utility equals u while every other choice gives at most u:

$$\begin{aligned}\pi_{\ell k}^{\text{all}} &= \int_0^\infty \Phi_\ell \, \partial_\ell G_k(\Phi) \, u^{-\theta_k^{\text{all}}-1} \exp[-G_k(\Phi) u^{-\theta_k^{\text{all}}}] \, du \\ &= \frac{\Phi_\ell \, \partial_\ell G_k(\Phi)}{\theta_k^{\text{all}} G_k(\Phi)} = \frac{\partial \log V_k^{\text{all}}}{\partial \log \Phi_\ell}.\end{aligned} \tag{B.28}$$

At any nest,

$$\frac{\partial \log V_{vk}}{\partial \log V_{ak}} = \frac{\overline{B}_{va,k} V_{ak}^{\theta_v}}{\sum_{u \in \mathcal{C}(v)} \overline{B}_{vu,k} V_{uk}^{\theta_v}} = \pi_{a|v,k}. \tag{B.29}$$

The chain rule therefore gives the product of conditional shares in (B.25) and the preceding probability expression. For two choices a and d in the same nest, $\pi_{a|v,k}/\pi_{d|v,k} = (\overline{B}_{va,k}/\overline{B}_{vd,k})(V_{ak}/V_{dk})^{\theta_v}$. Hence θ_v is the elasticity of conditional odds with respect to relative value. Holding other values fixed, the own-share elasticity is $\theta_v(1 - \pi_{a|v,k})$. In particular, for low-skill individuals,

$$\frac{\partial \log N_L^{\text{reg}}}{\partial \log V_L^{\text{reg}}} = \mu_L(1 - b_L^{\text{reg}}), \tag{B.30}$$

and for high-skill individuals the corresponding elasticity within sector g is $\mu_H(1 - b_{gH}^{\text{reg}})$.

B.2.4 Baseline normalization and sectors

Let $\pi_{a|v,k}^0$ denote the model's baseline conditional probability. At the location-pair margin, it is predicted by the calibrated amenities, commuting costs, and baseline economic conditions:

$$\pi_{\ell|\mathcal{D},k}^0 = \frac{(\Phi_\ell^0)^\chi}{\sum_{m \in \mathcal{D}^0} (\Phi_m^0)^\chi}, \qquad \ell \in \mathcal{D}^0, \tag{B.31}$$

where $\mathcal{D}^0$ contains the choices available at baseline. Internal commuting probabilities come from fitted gravity flows, so an observed zero on a feasible link need not imply a zero probability. Broader conditional probabilities reproduce the corresponding calibrated population, sector, and participation totals. They are distinct from amenities and from preference parameters in levels.

At a broader nest, the preference parameters satisfy

$$\overline{B}_{va,k} = \pi_{a|v,k}^0 \left(\frac{V_{vk}^0}{V_{ak}^0}\right)^{\theta_v}. \tag{B.32}$$

At the location-pair margin, $\overline{B}_{va,k} = 1$ and this identity follows from (B.31). Writing $\widehat{V}_{ak} = V_{ak}/V^0_{ak}$ yields, on the baseline choice set,

$$\widehat{V}_{vk} = \left[\sum_a \pi^0_{a|v,k} \widehat{V}^{\theta_v}_{ak}\right]^{1/\theta_v}, \qquad \pi_{a|v,k} = \frac{\pi^0_{a|v,k} \widehat{V}^{\theta_v}_{ak}}{\sum_u \pi^0_{u|v,k} \widehat{V}^{\theta_v}_{uk}}. \tag{B.33}$$

In particular, when the feasible bilateral choice set is unchanged,

$$\pi^1_{\ell|\mathcal{D},k} = \frac{\pi^0_{\ell|\mathcal{D},k}(\Phi^1_\ell/\Phi^0_\ell)^\chi}{\sum_{m\in\mathcal{D}^0} \pi^0_{m|\mathcal{D},k}(\Phi^1_m/\Phi^0_m)^\chi}. \tag{B.34}$$

The fixed amenities and commuting costs cancel from the utility ratios and remain embodied in the baseline predicted probabilities. No additional bilateral preference parameter enters this update. Baseline counts and amenity calibration are described in Appendix C.6.1.

For area r, let E^0_{rgk} be employment in sector g, $E^0_{rk} = \sum_g E^0_{rgk}$, and U^0_{rk} nonemployment. The sector allocation of nonworkers is

$$a^0_{rgk} = E^0_{rgk}/E^0_{rk}, \qquad U^0_{igk} = a^0_{rgk} U^0_{ik}, \qquad i \in r. \tag{B.35}$$

All sectors at a nonworking residence have the same income, prices, and amenities. Summing over sectors recovers U^0_{ik}, and the baseline work probability within every sector is $\pi^0_{rk} = E^0_{rk}/(E^0_{rk}+U^0_{rk})$. With hats denoting ratios to baseline values of the indices in (B.8) and (B.9),

$$\widehat{\mathcal{S}}_{rgk} = \pi^0_{rk}\widehat{\widetilde{\mathcal{W}}}^{\eta/\psi}_{rgk} + (1-\pi^0_{rk})\widehat{\mathcal{U}}^{\eta/\chi}_{rgk}, \qquad \widehat{\mathcal{Q}}_{rk} = \sum_g a^0_{rgk}\widehat{\mathcal{S}}^{\nu/\eta}_{rgk},$$

$$\pi(g \mid r,k) = \frac{a^0_{rgk}\widehat{\mathcal{S}}^{\nu/\eta}_{rgk}}{\widehat{\mathcal{Q}}_{rk}}, \qquad \pi(W \mid r,g,k) = \frac{\pi^0_{rk}\widehat{\widetilde{\mathcal{W}}}^{\eta/\psi}_{rgk}}{\widehat{\mathcal{S}}_{rgk}}. \tag{B.36}$$

The parameter ν therefore governs sector choice, including the associated value of nonemployment; η governs work versus nonwork within a sector. An unconditional employed-sector share also reflects participation and location responses. When $\eta = \nu$, these two value indices combine into a single employment and sector index after aggregating the identical nonemployment choices. For low-skill individuals, the national choice is not divided across sectors. For high-skill individuals, the sector shares update in the same way with $\widehat{\mathcal{V}}_{gH}$, and each sector keeps its own fixed national option.

B.3 Disposable income and consumption demand

Disposable income is the wage or nonemployment income plus the residential dividend,

$$y_\ell = \begin{cases} w^g_{jk}, & \text{if working,} \\ y^U_k, & \text{if nonemployed,} \end{cases} + \mathbf{1}\{i \in \mathcal{J}\}D_k + \mathbf{1}\{i = o\}\overline{D}_{ok}. \tag{B.37}$$

Nonemployment income is externally financed. For CZ residents, this reduces to (12). Outside dividends are fixed. An outside resident employed inside the CZ receives that workplace's endogenous wage. A feasible consumption choice requires $y_\ell > 0$.

B.3.1 Expenditure shares and home and work shopping

A worker's service composite is $n = n_i^\rho n_j^{1-\rho}$. For fixed locations, the consumption problem is

$$\begin{aligned} \max_{h,n_i,n_j,x>0} \quad & \alpha_\ell \log h + \gamma_\ell \rho \log n_i + \gamma_\ell (1-\rho) \log n_j + \beta_\ell \log x, \\ \text{subject to} \quad & q_i^R h + P_i n_i + P_j n_j + x \le y_\ell. \end{aligned} \tag{B.38}$$

The amenity and taste terms are constant within this problem. The first-order conditions equate each good's expenditure to its utility exponent divided by the marginal utility of income. Summing these conditions and using the budget constraint shows that the marginal utility of income is $1/y_\ell$. Consequently,

$$\begin{aligned} q_i^R h_\ell &= \alpha_\ell y_\ell, & P_i n_i &= \rho \gamma_\ell y_\ell, \\ P_j n_j &= (1-\rho)\gamma_\ell y_\ell, & x_\ell &= \beta_\ell y_\ell. \end{aligned} \tag{B.39}$$

Housing demand is

$$h_\ell = \frac{\alpha_\ell y_\ell}{q_i^R}, \tag{B.40}$$

and spending on housing, services, and the tradable good is $\alpha_\ell y_\ell$, $\gamma_\ell y_\ell$, and $\beta_\ell y_\ell$. For CZ residents, $\alpha_\ell = \alpha_H$ and $\gamma_\ell = \gamma = (1-\alpha_H)s_N$, giving (13). A nonworker buys services only from home, so $P_i n_i = \gamma_\ell y_\ell$, with the same housing and tradable demands. Thus a dollar of additional disposable income raises service expenditure by γ_ℓ, allocated between home and work according to employment status.

B.3.2 Indirect utility and the expenditure-share normalization

Define $R_W = \rho^\rho (1-\rho)^{1-\rho}$ and $R_U = 1$, with endpoint values interpreted by continuity. Substitution of the demands into (B.1) gives

$$\begin{aligned} U_\ell^*(\omega) &= z_\ell(\omega) K_\ell \frac{\mathcal{B}_\ell y_\ell}{(q_i^R)^{\alpha_\ell} \mathcal{P}_\ell^{\gamma_\ell}}, & K_\ell &= \alpha_\ell^{\alpha_\ell} \gamma_\ell^{\gamma_\ell} \beta_\ell^{\beta_\ell} R_e^{\gamma_\ell}, \\ \mathcal{P}_\ell &= \begin{cases} P_i^\rho P_j^{1-\rho}, & e = W, \\ P_i, & e = U. \end{cases} \end{aligned} \tag{B.41}$$

The housing share, nonhousing shares, and shopping allocation are fixed within each entry comparison. The constant K_ℓ depends only on area r and employment status e, so write it as K_{re}. It is common across location pairs within each such group, although it can differ between workers

and nonworkers and across areas. Homogeneity of the working and nonworking indices allows these constants to be absorbed into the participation shifters:

$$\overline{B}^{E}_{rgk} = \overline{B}^{E,\text{raw}}_{rgk} K^{\eta}_{rW}, \qquad \overline{B}^{U}_{rgk} = \overline{B}^{U,\text{raw}}_{rgk} K^{\eta}_{rU}. \tag{B.42}$$

Thus the normalized systematic value in (4) produces the same probabilities and inclusive values as the raw indirect utility when the participation shifters are expressed in the corresponding units. This fixed taste normalization preserves the original Cobb–Douglas preferences and expenditure shares without a bilateral preference parameter. Each shopping counterfactual exercise uses the same normalization in its entry and no-entry equilibria.

B.4 Population, employment, and labor supply

Multiplying the conditional shares yields a regional choice probability $\pi^{\text{all}}_{\ell k} = b^{\text{reg}}_{k} \pi(r \mid \text{reg}, k)\pi(g \mid r, k)\pi(e \mid r, g, k)\pi_{\ell|r,g,e,k}$. For outside-area workers, the last probability includes the job-location choice. The national share is $1 - b^{\text{reg}}_{k}$. Conditional shares sum to one at every margin, so $\sum_{\ell} \pi^{\text{all}}_{\ell k} = 1$ and $n_{\ell k} = \mathcal{N}_k \pi^{\text{all}}_{\ell k}$ exhausts the fixed national population.

Let F_{ijgk} denote working counts and U_{igk} nonworking counts by sector. The relevant employment and population totals are

$$\begin{aligned} U_{ik} &= \sum_{g} U_{igk}, & L^{g}_{jk} &= \sum_{i} F_{ijgk}, \\ N_k &= \sum_{i \in \mathcal{J}} \left(\sum_{j,g} F_{ijgk} + U_{ik} \right), & N &= \sum_{k} N_k. \end{aligned} \tag{B.43}$$

Workplace employment includes inbound commuters; resident employment includes outbound commuters. These sums distinguish changes in local jobs from changes in residents' employment. In particular, entry can raise workplace employment through commuting without changing the number of employed CZ residents.

B.5 Individual values and choice changes

Let $y_{\ell k}$ be disposable income in choice ℓ and $q_i = (1+t) r^{R}_{i}$ the residential rent for a local resident. Housing has the Cobb–Douglas expenditure share α_ℓ. For an employed individual living in i and working in j, the change in the systematic utility is

$$d_{\ell k} = \log \frac{y_{\ell k}}{y^{0}_{\ell k}} - \alpha_\ell \log \frac{q_i}{q^{0}_{i}} - (1 - \alpha_\ell) s_N \left[\rho \log \frac{P_i}{P^{0}_{i}} + (1-\rho) \log \frac{P_j}{P^{0}_{j}} \right]. \tag{B.44}$$

Income includes the current skill dividend alongside wages or nonemployment income. For a nonworker the shopping term is $\log(P_i / P^{0}_{i})$. Outside wages, nonemployment income, dividends, and prices are fixed. An outside resident employed inside the CZ receives the endogenous local wage.

Fixed amenities and commuting costs remain embodied in the baseline predicted probabilities.

For each bilateral choice, $d_{\ell k} = \log(\Phi_\ell / \Phi_\ell^0)$.

Set $v_{\ell k} = d_{\ell k}$. Each nest v has log inclusive-value change and conditional probabilities

$$v_{vk} = \frac{1}{\theta_v} \log \sum_{a \in \mathcal{C}(v)} \pi^0_{a|v,k} \exp(\theta_v v_{ak}), \tag{B.45}$$

$$\pi_{a|v,k} = \pi^0_{a|v,k} \exp\{\theta_v (v_{ak} - v_{vk})\}. \tag{B.46}$$

Multiplying the conditional probabilities along the nests containing a choice, and then by $\mathcal{N}_k$, gives its population. The baseline probabilities carry the calibrated amenities and commuting costs, so no additional amenity parameters enter. Unavailable choices have $v_{ak} = -\infty$. A nest with no feasible choices has zero value in levels, log value $-\infty$, and zero choice probability. A choice with nonpositive disposable income is unavailable. The baseline calibration requires positive income in every choice with positive baseline mass.

For two members u and z of the same nest v with positive shares in both equilibria,

$$\log \frac{\pi_{u|v,k}}{\pi_{z|v,k}} - \log \frac{\pi^0_{u|v,k}}{\pi^0_{z|v,k}} = \theta_v (v_{uk} - v_{zk}). \tag{B.47}$$

Thus a scale is a conditional log-odds response to the relative value. For a binary choice with share π, the elasticity of that share to its own value is $\theta_v(1-\pi)$. This distinction applies to participation and inward commuting as well as regional attachment. The two migration choices, between the region and the rest of the country and between residence inside and outside the CZ, share μ and therefore combine into a single nest.

B.6 Service demand and shopping price indices

The service bundle combines varieties from the two service sectors $g \in \mathcal{G}_N = \{N_H, N_L\}$. These sectors differ in skill intensity; each can employ both worker skills k. Appendix C.5.1 gives their sector classifications. Firms form a continuum, each supplying one differentiated variety under monopolistic competition with free entry. The mass m_{jg} is a continuous measure of the number of active firms in sector g and tract j, each charging p_{jg}. All varieties enter the same CES bundle with elasticity $\sigma_N > 1$, so the sector grouping introduces no additional consumption nest. Sector prices and establishment masses are determined separately by production costs and free entry.

Sector-specific shopping demand. Shopping and commuting share the per-minute travel-cost parameter κ. The shopping preference parameters satisfy $\zeta_{ijg} = \zeta_{ij}\xi_{jg}$. The bilateral shifter ζ_{ij} is common to both industries, while ξ_{jg} measures destination–industry appeal. Thus $T_{ijg} = \zeta_{ij}\xi_{jg} \exp\{-(\sigma_N - 1)\kappa t_{ij}\}$.

Each delivered variety price is $p_{jg} \exp(\kappa t_{ij})$; summing the dual price terms over varieties gives $T_{ijg} m_{jg} (p_{jg})^{1-\sigma_N}$, while fixed outside opportunities contribute T_i^O. CES expenditure minimization

at origin i therefore gives

$$P_i = \left[T_i^O + \sum_{j \in \mathcal{J}} \sum_{g \in \mathcal{G}_N} T_{ijg} m_{jg} (p_{jg})^{1-\sigma_N} \right]^{1/(1-\sigma_N)}. \tag{B.48}$$

Shephard's lemma gives the expenditure share received by each industry and destination:

$$\frac{\partial \log P_i}{\partial \log p_{jg}} = \frac{T_{ijg} m_{jg} (p_{jg})^{1-\sigma_N}}{P_i^{1-\sigma_N}} = \pi_{ijg}^I. \tag{B.49}$$

Individuals can also purchase services outside the CZ. The term T_i^O summarizes the contribution of outside varieties, prices, shopping preferences, and travel costs to the service price index at origin i. These outside opportunities are held fixed across equilibria. The outside shopping share and the expenditure adding-up condition are

$$\pi_i^O = \frac{T_i^O}{P_i^{1-\sigma_N}}, \qquad \sum_{j \in \mathcal{J}} \sum_{g \in \mathcal{G}_N} \pi_{ijg}^I + \pi_i^O = 1. \tag{B.50}$$

Although T_i^O is fixed, the outside expenditure share responds to local prices and entry through P_i. Lower local prices or additional local varieties reduce the share spent outside. Expenditure $\mathcal{E}_i \pi_i^O$ accrues to businesses outside the CZ and does not enter local nontradable revenue.

A worker living in i and employed at j allocates $\rho \gamma_\ell y_{ijgk}$ to purchases from home and $(1-\rho)\gamma_\ell y_{ijgk}$ to purchases from work. Summing the first budget over F_{ijgk} and the second over F_{jigk} gives workers' spending from origin i. Nonemployment income and expenditure shares are identical across sectors at a given residence and skill, so nonworkers contribute $\gamma U_{ik}(y_k^U + D_k)$ after summing their counts over sectors. These three contributions give (14). The sums include feasible residence and workplace locations inside and outside the CZ. Inbound commuters contribute spending from work and outbound commuters contribute spending from home. Spending originating in the outside location accrues outside the CZ. Revenue at local destinations equals the spending directed to them from CZ origins.

Aggregation by tract. The main-text shopping equations combine industries within each tract. For a tract with positive variety mass, define total varieties and the CES price index per variety by

$$\begin{aligned} m_j &= \sum_{g \in \mathcal{G}_N} m_{jg}, \\ p_j^N &= \left[\sum_{g \in \mathcal{G}_N} \xi_{jg} \frac{m_{jg}}{m_j} (p_{jg})^{1-\sigma_N} \right]^{1/(1-\sigma_N)}. \end{aligned} \tag{B.51}$$

Let $T_{ij} = \zeta_{ij} \exp\{-(\sigma_N - 1)\kappa t_{ij}\}$. Since $T_{ijg} = T_{ij}\xi_{jg}$, these definitions imply

$$T_{ij} m_j (p_j^N)^{1-\sigma_N} = \sum_{g \in \mathcal{G}_N} T_{ijg} m_{jg} (p_{jg})^{1-\sigma_N}. \tag{B.52}$$

Substituting into (B.48) gives (15). A tract with no varieties contributes zero to the price-index sum. Expenditure shares and revenue aggregate as

$$\pi_{ij}^I = \sum_{g \in \mathcal{G}_N} \pi_{ijg}^I, \qquad R_j^N = \sum_{g \in \mathcal{G}_N} R_{jg}. \tag{B.53}$$

Summing expenditure shares gives (16). Total nontradable revenue in tract j is

$$R_j^N = \sum_{i \in \mathcal{J}} \mathcal{E}_i \pi_{ij}^I. \tag{B.54}$$

All service revenue comes from the spending of individuals in the model. The tract price index reflects both industry prices and the composition of available varieties; the industry prices and entry responses remain distinct.

Changes in shopping shares and revenue. Define destination attraction and the origin denominator in changes by

$$\Omega_{jg} = \widehat{m}_{jg} (\widehat{p}_{jg})^{1-\sigma_N}, \qquad D_i = \sum_j \sum_{g \in \mathcal{G}_N} \pi_{ijg}^{I,0} \Omega_{jg} + \pi_i^{O,0}, \tag{B.55}$$

$$\pi_{ijg}^I = \frac{\pi_{ijg}^{I,0} \Omega_{jg}}{D_i}, \qquad \pi_i^O = \frac{\pi_i^{O,0}}{D_i}, \qquad \widehat{P}_i = D_i^{1/(1-\sigma_N)}. \tag{B.56}$$

Individual NT budgets follow the Cobb–Douglas demands and are assigned to home and work origins using ρ. Nonworkers shop from home. Revenue in each industry and tract is

$$R_{jg} = \sum_{i \in \mathcal{J}} \mathcal{E}_i \pi_{ijg}^I, \tag{B.57}$$

where spending $\mathcal{E}_i$ includes each individual's income and dividends.

B.7 Nontradable production, costs, and entry

The model contains two service sectors, $g \in \mathcal{G}_N = \{N_H, N_L\}$. They differ in skill intensity, but each employs both worker skills k. Productivity, skill shares, and fixed input requirements may differ across tracts and sectors. Wages, prices, and entry are determined separately for each sector. The parameters ϵ_m and σ_N are common to both, and the quantification gives the two sectors the same cost shares.

B.7.1 Production and unit cost

A firm in sector $g \in \mathcal{G}_N$ and tract j produces output Y from labor $\mathcal{L}$, commercial space H, and intermediate inputs X, omitting firm subscripts,

$$Y = A_{jg}\, \mathcal{L}^{\,\alpha_L^g} H^{\,\alpha_C^g} X^{\,\alpha_X^g}, \qquad \alpha_L^g + \alpha_C^g + \alpha_X^g = 1. \tag{B.58}$$

Here A_{jg} is productivity and α_L^g, α_C^g, and α_X^g are the expenditure shares on labor, commercial space, and intermediates. If L_k is the firm's employment of skill k, the two skills form a Cobb–Douglas labor bundle with fixed expenditure shares λ_{jgk}. Minimizing the wage bill required to produce one unit of this bundle gives its unit cost:

$$\mathcal{L} = \prod_k L_k^{\lambda_{jgk}}, \qquad c_{jg}^L = \prod_k \left(\frac{w_{jk}^g}{\lambda_{jgk}} \right)^{\lambda_{jgk}}, \qquad \sum_k \lambda_{jgk} = 1. \tag{B.59}$$

A worker of skill k costs the firm the wage w_{jk}^g. Minimizing the cost of output gives

$$c_{jg} = \frac{(c_{jg}^L)^{\alpha_L^g} (q_j^C)^{\alpha_C^g} (p^X)^{\alpha_X^g}}{A_{jg} (\alpha_L^g)^{\alpha_L^g} (\alpha_C^g)^{\alpha_C^g} (\alpha_X^g)^{\alpha_X^g}}, \qquad p_{jg} = \frac{\sigma_N}{\sigma_N - 1}\, c_{jg}. \tag{B.60}$$

Here $q_j^C = (1+t) r_j^C$ and $p^X = 1$. The constant in the denominator follows from the Cobb–Douglas production normalization. Monopolistic competition under CES demand implies the constant markup shown above. Differentiating unit cost with respect to input prices shows that labor, commercial space, and intermediates receive the constant shares α_L^g, α_C^g, and α_X^g of total cost, whatever the factor prices. Variable production and fixed inputs use the same technology, so these shares apply to their combined expenditure C_{jg}.

B.7.2 Factor demand and free entry

Cost minimization within the Cobb–Douglas labor bundle gives $w_{jk}^g L_k = \lambda_{jgk} c_{jg}^L \mathcal{L}$. Shephard's lemma applied to c_{jg} gives the same constant shares. Aggregating variable and fixed costs across firms within sector g gives the input demands

$$L_{jk}^g = \frac{\alpha_L^g \lambda_{jgk} C_{jg}}{w_{jk}^g}, \qquad H_{jg}^C = \frac{\alpha_C^g C_{jg}}{q_j^C}, \qquad X_{jg} = \frac{\alpha_X^g C_{jg}}{p^X}. \tag{B.61}$$

At the CES markup, variable cost is $(\sigma_N - 1) R_{jg}/\sigma_N$, leaving revenue less variable costs of R_{jg}/σ_N. Each firm also incurs an annual fixed cost f_{jg}. With m_{jg} active firms, free entry drives profits to zero:

$$f_{jg} = \varphi_{jg} c_{jg} m_{jg}^{1/\epsilon_m - 1}, \qquad \frac{R_{jg}}{\sigma_N} = m_{jg} f_{jg}. \tag{B.62}$$

Fixed inputs use the same technology as output, so their quantity per firm is $f_{jg}/c_{jg} = \varphi_{jg} m_{jg}^{1/\epsilon_m - 1}$. The parameter φ_{jg} scales this requirement and is held fixed across equilibria. These inputs are fixed

with respect to an individual firm's output; their cost depends on factor prices and the number of firms, which each firm takes as given. The elasticity of fixed inputs per firm with respect to the number of firms is $1/\epsilon_m - 1$. Thus, when $\epsilon_m < 1$, additional entry raises each firm's fixed input requirement; when $\epsilon_m = 1$, that requirement is constant.

The zero-profit condition implies that revenue covers variable and fixed costs: $C_{jg} = R_{jg}$. Substitution gives firm entry and the employment response within each industry:

$$m_{jg} = \left(\frac{R_{jg}}{\sigma_N \varphi_{jg} c_{jg}} \right)^{\epsilon_m}, \qquad d \log L^g_{jk} = d \log R_{jg} - d \log w^g_{jk}. \tag{B.63}$$

Holding φ_{jg} fixed, ϵ_m is the elasticity of the number of firms with respect to revenue relative to unit cost. At given revenue and unit cost, the elasticity with respect to φ_{jg} is $-\epsilon_m$. Higher fixed input requirements therefore support fewer firms. Skill and cost shares are fixed, so the employment response combines additional revenue with the change in the wage. Payroll in sector g is $\sum_k w^g_{jk} L^g_{jk} = \alpha^g_L R_{jg}$. The calibration of cost shares and fixed input requirements is in Appendix C.6.3.

B.7.3 Aggregation across service sectors

The production and entry equations apply separately to each service sector. Total nontradable employment by skill and commercial-space demand in tract j are

$$L^N_{jk} = \sum_{g \in \mathcal{G}_N} L^g_{jk}, \qquad H^{C,N}_j = \sum_{g \in \mathcal{G}_N} H^C_{jg}. \tag{B.64}$$

These totals enter property-market clearing and the employment measures in Section 4.6. Total varieties are $m_j = \sum_{g \in \mathcal{G}_N} m_{jg}$ and total revenue is $R^N_j = \sum_{g \in \mathcal{G}_N} R_{jg}$. The shopping price p^N_j is the CES price index per variety in (B.51), whereas each p_{jg} is the price charged by firms in sector g. Without a sector index, nontradable employment, commercial-space demand, varieties, and revenue denote the corresponding tract totals. These sums preserve the separate costs and entry responses of the two service sectors.

B.8 Tradable production and the entrant

For a tradable sector $g \in \mathcal{G}_T$ with output Y_{jg}, minimizing cost with the Cobb–Douglas technology gives

$$c_{jg} = \frac{(q^C_j)^{\alpha^g_C} (p^X)^{\alpha^g_X}}{A_{jg}} \prod_k \left(w^g_{jk} \right)^{\alpha^g_L \lambda_{jgk}}. \tag{B.65}$$

Fixed expenditure-share constants are absorbed into A_{jg}, and the intermediate input price is normalized to $p^X = 1$. Differentiating total cost $c_{jg} Y_{jg}$ with respect to each input price gives the labor

demand in (24),

$$L^g_{jk} = \frac{\alpha^g_L \lambda_{jgk} c_{jg} Y_{jg}}{w^g_{jk}}, \tag{B.66}$$

and the remaining input demands,

$$H^C_{jg} = \frac{\alpha^g_C c_{jg} Y_{jg}}{q^C_j}, \qquad X_{jg} = \frac{\alpha^g_X c_{jg} Y_{jg}}{p^X}. \tag{B.67}$$

The skill shares satisfy $\sum_k \lambda_{jgk} = 1$. Because output sells at the national price, the zero-profit condition (23) ties the skill-weighted wage a producer can pay to its productivity and commercial rent,

$$\alpha^g_L \sum_k \lambda_{jgk}\, d\log w^g_{jk} = d\log A_{jg} - \alpha^g_C\, d\log q^C_j. \tag{B.68}$$

Output Y_{jg} then adjusts until employment equals the labor supply at these wages, so a producer shrinks when its workers are offered higher wages elsewhere. Profits are zero in every location and sector.

Entry raises the productivity $A_{j_E g_E}$ of the entrant's location-sector by the factor that raises its employment by J^E in the entry equilibrium. The entrant is not distinguished from incumbent producers in its location-sector. It pays the sector–skill wages in its tract, and its demands for labor, commercial space, and intermediate inputs follow (B.66) and (B.67), so it competes with other producers for local commercial space.

B.9 Property supply, residential dividends, and property taxes

Residential and commercial supply follow (26), independently at each tract. Residential and commercial rents paid by users include the property tax: $q^R_i = (1+t) r^R_i$ and $q^C_j = (1+t) r^C_j$. Housing demand sums the individual demand $h_\ell = \alpha_\ell y_\ell / q^R_i$. Commercial demand sums the tradable and nontradable demands above. Equating each demand to supply gives (27).

Residential rental income $\mathcal{R} = \sum_i r^R_i H^R_i$ is distributed using fixed skill ownership shares. Hence

$$D_k = \frac{\omega^0_k \mathcal{R}}{N_k}, \qquad \sum_k N_k D_k = \mathcal{R}, \qquad d\log D_k = d\log \mathcal{R} - d\log N_k. \tag{B.69}$$

The differential applies to positive ownership shares. Dividends accrue to current residents, including renters and new arrivals; commercial rents accrue to absentee landlords.

The tax applies to pretax rent on all residential and commercial floor space, including the space used by the entrant's location-sector. With commercial rental income $\mathcal{R}^C$,

$$\mathcal{R}^C = \sum_{j \in \mathcal{J}} r^C_j H^C_j, \qquad \mathcal{T} = t\left(\mathcal{R} + \mathcal{R}^C\right). \tag{B.70}$$

Tax revenue leaves the modeled economy. It is not rebated to residents or spent locally. Nonemployment income comes from outside the CZ. The award is a payment to the entrant that does not change its hiring, and the model does not solve how it is financed.

B.10 Equilibrium accounting and feasible choices

For each entrant specification, equilibrium determines

$$\begin{aligned}
&\{w^g_{jk}, r^R_i, r^C_j, p_{jg}, m_{jg}\}_{i,j,k,g},\\
&\{F_{ijgk}, U_{ik}, n_{\ell k}, N^{\text{reg}}_k, N^{\text{nat}}_k, N_k, N\},\\
&\{L^g_{jk}, H^C_{jg}, R_{jg}\},\\
&\{h_\ell, P_i, \pi^I_{ijg}, \pi^O_i, \mathcal{E}_i, D_k\}.
\end{aligned}$$

These quantities satisfy individual choice and consumption, firms' input and entry conditions, labor and property-market clearing, the revenue account, finite-population accounting, and the residential portfolio described in Section 4.4.

In particular, the sector-by-sector labor-market conditions underlying (32) are

$$\sum_i F_{ijgk} = L^g_{jk}, \qquad g \in \mathcal{G},\ j \in \mathcal{J},\ k \in \{H, L\}. \tag{B.71}$$

The tradable conditions include the entrant's location-sector, whose labor demand reflects its higher productivity. The service conditions hold separately for each $g \in \mathcal{G}_N$ and each skill k. Summing commercial-space demands over all four sectors gives (27).

Modeled service destinations are tracts with positive pre-entry nontradable employment. Workplace–sector choices use the observed baseline employment support. The entrant's location-sector has positive baseline employment in every event, so entry opens no new choices. Primitives are calibrated jointly to pre-entry prices, employment, costs, spending, and budgets. Each skill has a positive current resident population, feasible consumption choices, and a housing expenditure share strictly between zero and one. Individual disposable income is positive for every available choice.

B.11 Place-based welfare and employment accounting

For a place a, either a county or the CZ, let V^s_{ak} be the nested inclusive value over choices with residence in a. Amenities, tastes, and preference parameters are common to the entry and no-entry calculations. Let Λ^0_ℓ collect the amenity, commuting, and price factors multiplying y_ℓ in (4), evaluated at no-entry prices. Equivalent variation EV_{ak} solves

$$\begin{aligned}
&V^0_{ak}(\text{EV}_{ak}) = V^1_{ak},\\
&\quad \Phi^0_\ell(x) = \mathbf{1}\{\ell \in \mathcal{A}^0_k\}\Lambda^0_\ell \max\{y^0_\ell + x, 0\}, \qquad i \in a.
\end{aligned} \tag{B.72}$$

Thus EV_{ak} is the uniform transfer, evaluated at no-entry prices and opportunities, that reproduces the entry inclusive value. Here $\mathcal{A}_k^0$ is the structurally available no-entry choice set. A transfer can restore an income-infeasible choice within this set. Place-level equivalent variation weights these skill-specific transfers by no-entry resident populations: $\mathrm{EV}_a = \sum_k N_{ak}^0 \mathrm{EV}_{ak}$ in total dollars and EV_a/N_a^0 per individual. The tables report these aggregates and both skill-specific transfers. This measure values residential opportunities in a place; individual welfare incidence can also be studied using matched transitions.

Workplace employment changes include the entrant's jobs, the adjustment at other tradable producers, and nontradable adjustment. The annual gross fiscal price for place a is $S^E/\Delta L_a$ when $\Delta L_a > 0$. The annual net fiscal measure uses $(\frac{1}{2}S^E - \Delta\mathcal{T})/\Delta L_a$, where the county pays half of the annual award S^E. The full commitment per net job is $\overline{S}^E/\Delta L_a$, where $\overline{S}^E$ is the full award commitment. The government account is CZ-wide, including when the employment denominator is a county.

B.12 Individual transitions and welfare

Equilibrium counts follow the nested shares in (B.25). Individual transitions use common taste draws across the no-entry ($s = 0$) and entry ($s = 1$) equilibria. Let $\mathcal{A}_k^s$ contain the choices with feasible consumption under specification s. Chosen locations and employment, and maximized indirect utility, are

$$\ell_k^s(\omega) = \arg\max_{a \in \mathcal{A}_k^s} z_a(\omega)\Phi_a^s, \qquad v_k^s(\omega) = \max_{a \in \mathcal{A}_k^s} z_a(\omega)\Phi_a^s. \tag{B.73}$$

The taste distribution and potential location–industry choices are fixed across counterfactuals. Feasibility responds to equilibrium prices and income.

The skill-k individuals moving from choice a to choice b number

$$T_{ab,k} = \mathcal{N}_k \Pr\{\ell_k^0 = a,\ \ell_k^1 = b\}, \qquad \sum_b T_{ab,k} = n_{ak}^0, \qquad \sum_a T_{ab,k} = n_{bk}^1. \tag{B.74}$$

The joint probabilities use the same draws in both equilibria. Summing them by sector, workplace, residence, or nonemployment gives movements between these categories. Employment changes are net stock changes; tradable worker transitions are measured at the workplace level.

Welfare comparisons use $v_k^0(\omega)$ and $v_k^1(\omega)$ for the same individuals. Each utility includes the dividend associated with the chosen current residence. Welfare averages among current residents also reflect the changes in composition generated by migration and sorting.

C Quantification and Equilibrium Computation

All monetary flows are annual and expressed in constant 2023 dollars. Superscript 0 denotes the pre-entry equilibrium. A subscript c identifies CZs in the calibration; it is omitted from the single-

CZ model in Section 4 and Appendix B.

This appendix supports Section 5 in the order of its subsections. Appendix C.1 collects the parameters, Appendices C.2–C.4 support the individual choices and budgets in Section 5.1, Appendix C.5 supports Section 5.2, and Appendix C.6 supports the construction of the pre-entry economy in Section 5.3. Appendix C.7 describes the equilibrium computation, and Appendix C.8 the exposure used in the validation of Section 5.4.

C.1 Common and CZ-specific parameters

Table C.1: Parameters used in the equilibrium calculations

Object	Value	Basis
Behavioral parameters		
Migration, (μ_H, μ_L)	$(2.32, 1.51)$	Estimated; 2010–2018 shift-share 2SLS by skill (Appendix C.2.3)
Sector choice, ν	1.64	Estimated; 2012–2019 QWI shift-share 2SLS (Appendix C.2.1)
Participation, η	3.32	Recruitment calibration to Fact 1 hiring sources
Inbound-to-pair ratios, (q_S, q_M, q_L)	$(0.67, 0.81, 1.00)$	Recruitment calibration; small/medium/large CZs, $\psi_e = q_{\text{size}(e)}\chi_{c(e)}$
Residence–workplace choice, χ	4.93–11.37	CZ-by-year LODES PPML; five-minute within-tract time
Service substitution, $\sigma_{N,c}$	3.95–18.37	Pooled CZ shopping PPML; $\kappa = 0.01$
Establishment entry, ϵ_m	0.40	County CBP regression, 2002–2022; descriptive calibration (Appendix C.5.3)
Accounting and individual calibration		
Service cost shares (L, C, X)	(0.53, 0.13, 0.34)	2012 BEA intermediates; 80/20 labor–space split
High-skill tradable cost shares	(0.49, 0.12, 0.39)	Same; NAICS 51, 52, 54, 55
Low-skill tradable cost shares	(0.27, 0.07, 0.67)	Same; NAICS 31–33
Housing-value shares, ω_{ck}^0	By event	ACS owner values; split equally between partners
Nonemployment income, y_k^U	By event and skill	ACS transfers among nonworkers
Housing share, $\alpha_{H,c}$	0.17–0.28	CZ household-median expenditure share
NT share of nonhousing spending, $s_{N,c}$	0.63–0.89	By event baseline, 10th–90th percentile; service revenue inversion (BEA benchmark 0.48)
Maintained spending, travel, technology, and policy parameters		
Home-origin shopping share, ρ	0.74	74/26 home/work routing share
Travel cost, κ	0.01 per minute	Common to commuting and shopping
Residential supply, $\epsilon_{R,i}$	Tract-specific	Mapped Baum-Snow–Han values; eight-neighbor extension
Commercial supply, ϵ_C	0.60	Maintained commercial-space benchmark
Baseline property tax, t_c^0	0.28	Rate adjusts after entry to finance service needs
Resident-service share of county spending, ς	0.79	2012 Census of Governments
Entrant jobs, industry, site, added space	By event	Baseline capacity rule; 50% tax exemption on new space

Notes: Migration differs by skill; sector choice and participation are common across skills. Border scales vary by event through three CZ-size ratios; commuting and shopping scales also vary by CZ. The recruitment calibration uses all 55 original events (Section 5.1). Cost shares list labor, commercial space, and purchased intermediate inputs (Appendix C.5.2). Appendix C gives the estimands and diagnostics. The separate local-benefit exercise replaces the common 50-percent exemption with event-specific fractions for 54 scenarios (Appendix F.1); its labor parameters remain fixed.

Table C.1 reports the parameters. The tax, travel-cost, technology, establishment-entry, and commercial-supply parameters are common across CZs. Residential supply elasticities vary by tract. Housing expenditure shares, shopping elasticities, ownership shares, and location-specific shifters

are calibrated using the corresponding CZ data. The model uses migration scales $\mu_H = 2.32$ and $\mu_L = 1.51$, sector choice $\nu = 1.64$, and participation $\eta = 3.32$. The boundary scale varies across event baselines according to $\psi_e = q_{\text{size}(e)}\chi_{c(e)}$, where size$(e)$ is the event's baseline CZ-size group and $\chi_{c(e)}$ is the fixed residence–workplace scale of the event's CZ. The ratios for small, medium, and large CZs are $(0.67, 0.81, 1.00)$. Both worker skills face the same boundary scale within an event. Baseline wages, nonemployment rates, and preference parameters may differ by skill.

The nontradable basket comprises private retail, education, health, accommodation and food, and other services (NAICS 44–45, 61, 62, 72, and 81). Appendix C.5.2 derives the production cost shares of the two tradable groups and the nontradable composite, and Appendix C.6.3 describes nontradable fixed input requirements.

C.2 Sector choice, migration, and recruitment

This subsection supports the sector-choice, migration, and participation parameters in Section 5.1.

C.2.1 A public benchmark for the sector-choice elasticity

I discipline the common sector-choice elasticity, ν, using public Quarterly Workforce Indicators (QWI) employment and earnings data for 2012–2019. The preferred pooled employment-stock response is 1.64. The estimation compares tradable and nontradable employment, pooling the model's high- and low-skill industry categories within each broad sector. It provides an external benchmark for the common sector scale; its structural interpretation requires the restrictions discussed below.

Data and specification. Employment is full-quarter jobs (`EmpS`) and earnings are average monthly earnings of full-quarter employees (`EarnS`), for private-sector workers aged 25 and older. Nontradables comprise NAICS 44–45, 61, 62, 72, and 81; tradables comprise the remaining covered private industries. I use published broad-industry totals and sum employment and earnings bills across counties and industries within each CZ, education group, and sector before calculating average earnings.[35]

For employment L_{cks} and average annualized earnings w_{cks} in CZ c, education group k, and sector $s \in \{T, N\}$, I estimate

$$\Delta \log L_{cks} = \beta^{\text{QWI}} \Delta \log w_{cks} + \alpha_{ck} + \delta_{d(c)s} + \gamma_{ks} + \varepsilon_{cks}, \tag{C.1}$$

where Δ denotes the 2012–2019 change and $d(c)$ is the Census division containing CZ c. The fixed effects absorb CZ-by-education, division-by-sector, and education-by-sector differences in growth. Education distinguishes workers with and without a bachelor's degree; the pooled specification imposes a common coefficient. Each sector receives half its CZ–education group's fixed 2012 bench-

[35] A county–industry–education–year observation is retained when at least three quarters have positive, finite employment and earnings with accepted QWI release codes (1, 6, or 10). Annual employment averages those quarters, and the annualized earnings bill averages $12 \times$ `EmpS` $\times$ `EarnS`. Suppressed observations are not imputed.

mark employment weight, held constant across specifications.[36] With equal sector weights within each CZ–education group, the coefficient also equals the IV estimate from the relative-outcome regression

$$\Delta \log\left(\frac{L_{ckN}}{L_{ckT}}\right) = \beta^{\text{QWI}} \Delta \log\left(\frac{w_{ckN}}{w_{ckT}}\right) + \widetilde{\delta}_{d(c)} + \widetilde{\gamma}_k + \varepsilon_{ckN} - \varepsilon_{ckT}, \tag{C.2}$$

where $\widetilde{\delta}_d = \delta_{dN} - \delta_{dT}$ and $\widetilde{\gamma}_k = \gamma_{kN} - \gamma_{kT}$.

Instrument and identification. Because earnings and employment are jointly determined, I estimate equation (C.1) by weighted two-stage least squares. The instrument combines pre-period industry employment shares with subsequent employment growth outside the focal CZ:

$$a_{cksi} = \frac{\overline{L}_{cki}}{\sum_{j \in \mathcal{I}_s} \overline{L}_{ckj}}, \qquad Z_{cks} = \sum_{i \in \mathcal{I}_s} a_{cksi} G_{i,-c}, \tag{C.3}$$

$$G_{i,-c} = \log\left(\frac{L_{i,-c,2019}}{L_{i,-c,2012}}\right), \tag{C.4}$$

where $\overline{L}_{cki}$ is mean industry employment over 2007–2011, $\mathcal{I}_s$ contains the broad industries in sector s, and $L_{i,-c,t}$ pools employment across education groups outside the focal CZ. External growth uses the fixed source sample of 2,365 counties in 30 states, yielding 19 broad-industry shocks. Both stages use the same fixed effects and weights; the instrument for equation (C.2) is $Z_{ckN} - Z_{ckT}$.

Identification requires instrument relevance and orthogonality between residualized exposure and unobserved changes in relative labor supply, conditional on the fixed effects. For a labor-supply interpretation, external growth must affect relative employment through earnings rather than through correlated changes in sector amenities, nonwage compensation, or worker composition, and average earnings must track the opportunities of comparable workers. Predetermined shares and the leave-out construction do not establish these restrictions by themselves (Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022, 2025). Standard errors are clustered by CZ. Anderson–Rubin confidence intervals address weak instruments under this covariance assumption; they do not account for dependence across CZs arising from shared industry exposures (Adão et al., 2019).

Estimates. The preferred pooled estimate is $\widehat{\beta}^{\text{QWI}} = 1.64$, with a CZ-clustered standard error of 0.51, first-stage $F = 15.03$, and an Anderson–Rubin 95 percent confidence interval of $[0.79, 3.24]$. Table C.2 compares specifications using identical observations, instruments, and weights. Removing regional sector controls yields 0.80 (SE 0.37; $F = 36.68$). Additionally replacing CZ-by-education effects with plain CZ effects yields 1.16 (SE 0.32; $F = 37.05$) and allows differences in growth between education groups within a CZ to contribute to identification. The estimates therefore remain sensitive to the treatment of regional and education-specific growth.

[36] The benchmark weight is the 2012 employment total in NAICS 31–33, 44–45, 51–55, 61, 62, 72, and 81. It is not recalculated when the covered tradable industries are expanded.

Table C.2: Employment-stock responses across tradable and nontradable industries

Specification	Estimate	SE	First-stage F	AR 95% interval
Preferred controls	1.64^{***}	0.51	15.03	$[0.79, 3.24]$
No regional sector effects	0.80^{**}	0.37	36.68	$[0.05, 1.60]$
CZ and education-by-sector effects	1.16^{***}	0.32	37.05	$[0.57, 1.90]$

Notes: All rows pool education groups and use 2012–2019 changes, 2007–2011 exposure shares, and identical observations, instruments, and weights. Preferred controls are CZ-by-education, division-by-sector, and education-by-sector fixed effects. The second row removes division-by-sector effects; the third also replaces CZ-by-education effects with CZ effects. Standard errors are clustered by CZ. The first-stage statistic is the squared cluster-robust t statistic on the excluded instrument. AR intervals invert the Anderson–Rubin test using the same clustering and an $F(1, C-1)$ reference distribution, where C is the number of CZs. Stars indicate $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$ for two-sided CZ-clustered Wald tests of a zero coefficient, using t_{C-1} with $C = 378$; they do not use the AR tests.

Interpretation and the four-sector model. I use the pooled estimate as an external benchmark for the common sector scale $\nu = 1.64$ in Appendix C.2.2. The regression measures an employment-stock response that can include migration, commuting, participation, and composition. In the four-sector model, ν governs choice among sector opportunity values. Mapping the estimated coefficient into ν therefore requires restrictions connecting those values to observed earnings and accounting for the other adjustment margins. A structural calibration could instead target the model's analogue of equation (C.1), computed with the same aggregation, weights, and horizon.

The four-sector extension yields 0.80 (SE 0.36; first-stage $F = 18.50$; AR 95 percent interval $[0.13, 1.77]$). It uses more detailed industry data and a slightly smaller complete-choice sample, so the comparison also changes data coverage. Pairwise diagnostics vary substantially: the N_H–N_L estimate is -1.46 (SE 0.40; $F = 9.88$), with an AR interval of $[-3.02, -0.80]$. These findings qualify the interpretation of a common elasticity across all four choices. Dynamic worker-transition estimates, such as Artuç et al. (2010) and Caliendo et al. (2019), concern sector transitions governed by moving costs and continuation values. They provide context for sector mobility, but their parameters are not directly comparable to this seven-year employment-stock response.

C.2.2 Recruitment calibration

The recruitment calibration sets only the participation elasticity and the size-group border ratios reported in Table C.1, holding sector choice and migration at their estimates. Section 5.1 describes the targets, the twelve calibration events, and the fit.

C.2.3 Migration elasticity

This subsection derives the shift-share regression behind the migration elasticities in Section 5.1. For high-skill individuals, who choose a sector before a region, it treats the CZ's skill-level population as one regional choice, which approximates the within-sector margin when sector shares change little.

From the model to the estimating equation. Equation (B.11) gives the regional share of skill-k individuals. Its log odds are

$$\log \frac{b_k^{\text{reg}}}{1 - b_k^{\text{reg}}} = \log \frac{\overline{B}_k^{\text{reg}}}{\overline{B}_k^{\text{nat}}} + \mu_k \left(\log V_k^{\text{reg}} - \log V_k^{\text{nat}} \right). \tag{C.5}$$

Treating each commuting zone c as a region and the rest of the country as the national choice, the odds are $P_{ck}/(P_k^{US} - P_{ck})$. Differencing between 2010 and 2018 removes the preference parameters, and the change in the national value, which is common to all CZs within a skill, is absorbed by the division-by-skill effects:

$$\Delta \log \frac{P_{ck}}{P_k^{US} - P_{ck}} = \delta_{d(c),k} + \mu_k \, \Delta \log V_{ck}^{\text{reg}} + u_{ck}, \tag{C.6}$$

where u_{ck} collects changes in preference parameters and amenities relative to the division. Because the outcome is in log odds, its coefficient is μ_k itself; the population elasticity in equation (B.30) multiplies it by $1 - b_k^{\text{reg}}$, which is close to one for a single CZ. The horizon is one eight-year difference, and the coefficient is not annualized.

By equation (B.3), the change in an inclusive value is, to first order, the share-weighted average of the changes in the values below it, so $\Delta \log V_{ck}^{\text{reg}}$ averages $\Delta \log \Phi_\ell$ over the CZ's choices. From equation (4), holding amenities and nonwage income fixed, the value of a working choice changes by

$$\Delta \log \Phi_\ell = \frac{w_{jk}^g}{y_\ell} \, \Delta \log w_{jk}^g - \alpha_{H,c(i)} \Delta \log q_i^R - (1 - \alpha_{H,c(i)}) s_N \, \Delta \log \mathcal{P}_\ell. \tag{C.7}$$

The estimating equation is

$$\Delta \log \frac{P_{ck}}{P_k^{US} - P_{ck}} = \delta_{d(c),k} + \mu_k \Delta \mathcal{V}_{ck} + u_{ck}, \qquad \mathcal{V}_{ck} = \log(e_{ck} w_{ck}) - 0.25 \log r_c. \tag{C.8}$$

Its value proxy replaces the average of these changes with the change in expected earnings, $e_{ckt} w_{ckt}$, less 0.25 times the change in rent. The coefficient on $\Delta \mathcal{V}_{ckt}$ therefore recovers μ_k when the proxy's weights match the model's: a unit weight on earnings against w_{jk}^g / y_ℓ, and 0.25 on rent against $\alpha_{H,c(i)}$. Changes in service prices, dividends and nonemployment income enter the error. The sample also covers individuals aged 25–64 rather than everyone aged 18 or older.

Instrument and inference. Local earnings respond to local population through u_{ck}, so I instrument $\Delta \mathcal{V}_{ckt}$ with $Z_{ck} = \sum_n \omega_{ckn} g_{n,-c}$, where ω_{ckn} is the 2010 share of skill-k employment in commuting zone c in industry n and $g_{n,-c}$ is national CBP employment growth in industry n outside the CZ, with industries harmonized across classification vintages. The identifying assumption is $E[Z_{ck} u_{ck} \mid d(c)] = 0$: national industry growth outside the CZ affects local population only through the value of living there, conditional on division effects. Population growth can also reflect aging, mortality, and education transitions, which would violate this assumption if they were correlated

with the industry mix. I estimate equation (C.8) by two-stage least squares with equal weights across commuting zones, using the 2010 and 2018 ACS five-year files and QWI earnings by skill. CZs with similar industry mixes are exposed to the same shocks, so Table C.3 reports standard errors clustered by three-digit industry shock as well as by commuting zone (Adão et al., 2019; Borusyak et al., 2022). The low-skill first stage is below ten, and both confidence intervals are wide. A matched 2002–2010 estimate is not available in comparable data.

Table C.3: Migration elasticity estimates

	High skill	Low skill
Estimate, μ_k	2.32^{***}	1.51^{***}
Standard error, clustered by CZ	0.75	0.51
Standard error, clustered by three-digit industry shock	1.13	1.24
First-stage Wald statistic	10.71	8.41
95% confidence interval	[1.35, 5.60]	[0.90, 4.40]
Commuting zones	558	559

Notes: Stars use the first reported standard errors, clustered by CZ, and two-sided t_{C-1} tests of a zero coefficient, where C is the skill-specific CZ count. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$. The industry-shock clustered standard errors and weak-instrument diagnostics are reported separately.

C.3 Commuting and shopping

This subsection supports the commuting and shopping parameters in Section 5.1.

C.3.1 Commuting

The commuting scale χ and the cost per minute κ are common to both skills. Within a lowest nest $\mathcal{B}$, equation (B.6) gives the share of the residence–workplace pair $\ell = (i, j)$ as Φ_ℓ^χ relative to the nest total, with no pair-specific preference parameter, and the conditional shares in equation (C.29) reproduce the baseline allocation. The baseline flows of skill k are the commuting gravity weights balanced to the observed residence and workplace margins (Appendix C.6.1),

$$F_{ijk}^0 = a_{ik}\, b_{jk} \exp(-\chi\kappa\, t_{ij}), \tag{C.9}$$

where a_{ik} and b_{jk} are the balancing factors that match the margins. Cobb–Douglas pair values factor into residence and workplace components as shown in Appendix C.6.2. These components determine the economic interpretation of the balancing factors. Changes in income and prices update the fitted shares through $\hat{\Phi}_\ell^\chi$ (Appendix B.5). Equation (C.9) is a gravity equation, and the corresponding estimating equation is

$$E[F_{ij}] = \exp\left(\alpha_i^C + \delta_j^C + \beta_c^C t_{ij}\right), \qquad \beta_c^C = -\chi\kappa, \tag{C.10}$$

where the residence and workplace effects α_i^C and δ_j^C take the place of the balancing factors.

I estimate equation (C.10) by PPML separately for each commuting zone, which accommodates the many pairs with no commuters (Santos Silva and Tenreyro, 2006; Dingel and Tintelnot, 2026), using 2019 LODES residence–workplace flows of all workers, since the gravity weights are common to both skills, and 2020 driving times on routes of up to 180 minutes. The coefficient is identified from variation in travel time within residences and workplaces, and $\chi_c = -\beta_c^C/0.01$ under the normalization $\kappa = 0.01$ per minute. The scale χ governs where workers live and work within the CZ and is distinct from ψ_e, which governs whether residents of adjacent CZs work inside it.

C.3.2 Shopping

Using the industry aggregation in (B.52), spending from origin i at destination j in (16) is

$$\mathcal{E}_i \pi_{ij}^I = \underbrace{\mathcal{E}_i P_i^{\sigma_N - 1}}_{\text{origin}} \underbrace{\sum_{g \in \mathcal{G}_N} \xi_{jg} m_{jg} \left(p_{jg}\right)^{1-\sigma_N}}_{\text{destination}} \zeta_{ij} \exp\{-(\sigma_N - 1)\kappa t_{ij}\}. \tag{C.11}$$

Visits V_{ijc} measure trips rather than dollars. If spending per visit varies only by origin or by destination, the fixed effects absorb it. Writing the taste weight as $\zeta_{ij} = \exp\{\gamma_c \mathbf{1}[i = j] + \xi_c \mathbf{1}[\text{same county}]\}$ gives the estimating equation

$$E[V_{ijc} \mid \text{controls}] = \exp\{\alpha_{ic} + \delta_{jc} + \beta_c^S t_{ij} + \gamma_c \mathbf{1}[i = j] + \xi_c \mathbf{1}[\text{same county}]\}, \tag{C.12}$$

with $\beta_c^S = -(\sigma_{N,c} - 1)\kappa$.

I estimate equation (C.12) by PPML separately for each commuting zone, using 2019 visits and 2020 routes. The same-county term is included where supported, standard errors are clustered by origin and destination, and the routed support includes positive travel times up to 180 minutes, with five minutes assigned to trips within a tract. Visits are pooled over NAICS 44, 45, 61 excluding 6113, 722, and 81; education visits proxy for the health-care network and are counted once. The standard error of $\sigma_{N,c}$ is the coefficient's standard error divided by 0.01.

Origin and destination effects absorb how attractive each location is, but bilateral tastes and trip chaining can still covary with travel time. Within-CZ decay also does not identify how much spending leaks across commuting-zone boundaries, and baseline destination shares reflect the revenue balancing of Section 5.3 as well as the estimated decay.

From the coefficient to the elasticity. A pooled coefficient for each CZ's nontradable basket determines

$$\kappa = 0.01 \text{ per minute}, \qquad \sigma_{N,c} = 1 - \frac{\beta_c^S}{0.01}. \tag{C.13}$$

The coefficient κ is common to shopping and commuting and measures the exponential travel-cost increment per minute. Events in a CZ share $\sigma_{N,c}$, which governs substitution, markups, and free entry. The shifters T_{ij} and T_i^O incorporate shopping decay, and the destination productivities

recovered in Appendix C.6.4 align spending flows with baseline revenue. Visit data identify destinations; expenditure accounts supply spending totals, with retail expenditure measured as the merchant margin. Shopping weights are fixed across counterfactuals.

C.4 Individual budgets

This subsection supports the individual budgets in Section 5.1.

C.4.1 Income and the housing share

Workers receive wage income w^g_{jk}. Nonemployment income y^U_k is measured using public assistance, Supplemental Security Income, and Social Security payments among nonworkers. It is paid during nonemployment. Housing expenditure measures rental-equivalent services for renters and owners. Income and housing measures are per individual and use survey weights.

The population comprises individuals aged 18 or older. Skill-specific nonworker-to-worker ratios from pre-entry ACS microdata scale the observed resident worker counts. The accessible outside area pools neighboring CZs. Its population ratios for individuals aged 25–64 are adjusted to cover everyone aged 18 or older using available neighboring-state microdata; available-state age adjustments also bridge the national population stocks. These geographic and demographic approximations preserve a common finite-population decomposition.

Outside housing dividends divide observed rental-equivalent residential rental income, net of the property tax, according to outside ownership value shares and populations. Outside property tax revenue is not rebated. These dividends, outside residential rents, and the outside housing expenditure share are fixed in the counterfactuals. The local portfolio account reconciles individual income and rental income.

For each CZ, let $\widehat{E}_{H,c}$ be the weighted total pre-entry rental-equivalent housing-service expenditure and $\widehat{Y}^{\text{inc}}_c$ the total after-tax worker earnings and nonemployment income in the same sample. The residential portfolio distributes $\widehat{R}_c = \widehat{E}_{H,c}/(1+t_c)$ to the sample's adults. Sum income and the measured dividends within each observed household to obtain Y_{mc}. If E_m is its whole-household service bill and w_{mc} its survey weight times the PUMA-to-CZ fraction, the constant share and the implied service share are

$$\alpha_{H,c} = \text{wmed}_{m\in c}\left(\frac{E_m}{Y_{mc}}; w_{mc}\right), \qquad \gamma_c = s_N(1-\alpha_{H,c}). \tag{C.14}$$

Each observed household receives its survey weight once when calculating the median, and the model applies the resulting share to each individual's budget. The median targets a typical household's expenditure burden, and aggregate expenditure remains a separate diagnostic (Appendix A.7). Events sharing a CZ use the same earliest-pre-entry sample and parameter. The outside expenditure share uses the corresponding household median in its existing neighboring-CZ sample and fixed outside resource account. In the model's baseline, with total local income $Y^{\text{inc},0}_c$, producer housing

rental income satisfies

$$R_c^0 = \frac{\alpha_{H,c} Y_c^{\text{inc},0}}{1 + t_c - \alpha_{H,c}}. \tag{C.15}$$

The measured portfolio defines the empirical median's resource denominator; equation (C.15) determines rental income in the model's baseline. The portfolio divides this income according to ownership shares and current skill populations. Demanded housing at the observed residential rents anchors each tract's residential supply schedule. Constant expenditure shares approximate the housing Engel curve at the median household share, and the implied marginal service share is $s_N(1 - \alpha_{H,c})$. Appendix C.4.2 estimates that Engel curve.

C.4.2 The housing Engel curve

Housing expenditure determines how individual resources translate into demand for services and how residential rents enter location choices. The model uses a Cobb–Douglas housing share $\alpha_{H,c}$ for each CZ. The empirical calibration matches the household-weighted median housing-service expenditure share in the pre-entry sample. An auxiliary Engel-curve exercise describes the income gradient in housing shares and clarifies the distinction between a household budget share and a marginal expenditure parameter.

Estimating the Engel curve. The auxiliary exercise uses pre-entry ACS person and housing microdata for 15 CZs, selecting the earliest available pre-entry year within each CZ. Households are assigned fractionally from PUMAs to CZs; the crosswalk uses tract-population weights for 2010 PUMAs and area weights for 2000 PUMAs. All amounts are annual and expressed in 2023 dollars. Let n_m denote the number of adults in household m, ω_{mc} its household survey weight times the geographic allocation fraction, and $a_{mc} = n_m \omega_{mc}$ its adult weight. Per-adult resources in this auxiliary specification are reported household income less five percent of household wage income, divided by n_m. Positive-resource households below the within-CZ adult-weighted first percentile are trimmed. The fit uses owners and renters across the remaining income distribution, comprising 117,140 distinct households.

For renters, expenditure is annual contract rent. For owners, this particular auxiliary specification uses reported selected monthly owner costs when positive, with a rental-imputation fallback. Reported costs cover 99.97 percent of the included owner adult weight. The owner measure includes financing payments, including mortgage principal, and therefore differs from rental-equivalent service consumption. A survey-weighted renter regression of log rent on log rooms, bedrooms, structure type, and PUMA–year effects supplies the fallback and a common reference-dwelling price index q_m. Predicted rents are retransformed using a weighted smearing factor. Reference-dwelling rents are normalized to have an adult-weighted geometric mean of one.

Writing per-adult expenditure as e_m and resources as y_m, the auxiliary weighted least-squares

regression is

$$\frac{e_m}{y_m} = \alpha^{SG} + \delta_c \frac{q_m}{y_m} + u_{mc}, \qquad \delta_c = (1 - \alpha^{SG})\bar{h}_c. \tag{C.16}$$

The intercept is common across CZs and the price-to-income slopes are CZ-specific. The estimate is $\widehat{\alpha}^{SG} = 0.07$, with a conditional bootstrap standard error of 0.33 percentage points. The bootstrap resamples PUMA clusters; 195 of 199 draws yield usable estimates, holding the price and housing-imputation first stages fixed. This uncertainty measure omits first-stage and geographic-assignment uncertainty. The fitted intercept is the expenditure-share limit as q_m/y_m approaches zero, and hence an extrapolated limit rather than the observed average housing share. The corresponding expenditure relationship is

$$e_m = \alpha^{SG} y_m + (1 - \alpha^{SG}) q_m \bar{h}_c, \tag{C.17}$$

$$\left.\frac{\partial e_m}{\partial y_m}\right|_{q_m} = \alpha^{SG}, \qquad \frac{e_m}{y_m} = \alpha^{SG} + (1 - \alpha^{SG})\frac{q_m \bar{h}_c}{y_m}. \tag{C.18}$$

The derivative is structural only under the maintained demand and resource assumptions. Cross-sectional income variation can also reflect saving, measurement error, housing quality, financing arrangements, and household heterogeneity. Mean housing shares of approximately 24–31 percent in this auxiliary sample can therefore coexist with an estimated intercept near 6.70 percent.

Comparison with Tsivanidis. Tsivanidis (2026), Section 5.1, sets the long-run housing share $1-\beta = 0.24$ to match Bogotá's high-income expenditure pattern. His Appendix D.7 calibrates wages and the skill-productivity parameter to the wage premium, and then matches the GEIH average housing expenditure share of 0.31 using the housing minimum, jointly with the car-expenditure calibration. The reported post-period minimum is 1.21 in his model's units. My estimate of 0.07 is obtained from a pooled ACS regression with CZ-specific slopes. The estimators, housing measures, resource definitions, and populations differ; the estimates are not a controlled comparison of preferences between the United States and Bogotá.

Measurement and specification sensitivity are quantitatively material. A separate ACS exercise measures owner housing with rental equivalents, reconciles the residential portfolio with alternative replacements of combined investment income, and reserves households for validation. Its pooled full-sample intercepts range from 1.56 to 2.39 percent across the resource conventions, while held-out renter comparisons can favor the fixed 0.07 value. Owner service expenditures are imputed, so the combined fit does not independently validate owner consumption. These results identify sensitivity to measurement and accounting; they do not isolate a single explanation for the gap with Tsivanidis. The auxiliary Engel curve is evidence on housing demand, with its assumptions and limits reported alongside the estimate.

C.4.3 Housing-value shares by skill

Let V_a be the reported housing value of owner-occupied household a, v_a its survey weight, and f_{ak} the fraction attributed to skill k. Value is split equally between the householder and a spouse or partner when both are present; a householder without one receives the full attribution. The pre-entry ownership share is

$$\omega_{ck}^0 = \frac{\sum_{a \in c,\, \text{owner}} v_a V_a f_{ak}}{\sum_{a \in c,\, \text{owner}} v_a V_a}, \qquad \sum_k \omega_{ck}^0 = 1. \tag{C.19}$$

Pre-entry ACS microdata discipline these shares, with PUMA-to-CZ allocation providing the geographic mapping. Value shares incorporate group sizes, homeownership rates, and property values. The weights remain fixed across equilibria and allocate each skill's share of current residential rental income equally among its current residents.

C.5 Firms and space

This subsection supports Section 5.2.

C.5.1 Industry groups and baseline allocation

Tradable industry type is defined by the 2000 Census bachelor's-degree employment share at four-digit NAICS. A tradable industry is high skill when its share exceeds the unweighted median across industries, 0.21. High-skill nontradables comprise private education and health (NAICS 61 and 62); low-skill nontradables comprise private retail, accommodation and food, and other services (NAICS 44–45, 72, and 81). This industry label is distinct from individual education. The productive sectors are T_H, T_L, N_H, N_L, and both worker skills are available in each. The tradable basket comprises NAICS 31–33, 51, 52, 54, and 55; private nontradables comprise 44–45, 61, 62, 72, and 81.

LODES supplies tract employment totals over event years $-4, -3, -2$. Within each broad tradable industry, I pool Infogroup establishment employment over these years and apply its high- and low-skill industry shares to the corresponding LODES mean. All classified industries in NAICS 51, 52, 54, and 55 belong to T_H; manufacturing employment is split using its local industry mix. Missing or ambiguous manufacturing allocations use the classified manufacturing mix in the same tract, then the county, then the CZ. This rule applies to 2.89% of tradable employment across event baselines, and each allocation retains its source flag. Nontradable totals follow the fixed NAICS groups: private NAICS 61 and 62 enter N_H, and private NAICS 44–45, 72, and 81 enter N_L.

QWI private-sector employment and earnings provide the association between worker education and industry type. High-skill workers have a bachelor's degree or more; the remaining education categories form the low-skill group. QWI education categories describe workers aged 25 and above, so their industry association is used as a prior for the model's individual employment allocation. I retain industry-years with released information for both worker skills. Employment composition uses beginning-of-quarter employment; earnings use full-quarter employment to weight annualized

monthly mean earnings, expressed in 2023 dollars. The wage and employment samples retain their separate disclosure flags. County estimates pool the three pre-entry years. Where a county lacks usable information in any year for either worker skill, the corresponding CZ and event window supply the prior. This geographic fallback accounts for 0.08% of productive employment across event baselines. Coverage files retain suppressed and unclassified industries and compare released detailed employment with broad-industry controls.

Within a tract and broad sector $b \in \{T, N\}$, let L^b_{jH} and L^b_{jL} denote the two industry totals, so that broad sector b and industry type $f \in \{H, L\}$ together identify sector $g = b_f$, and let L^b_{jk} denote the existing totals by worker skill. Iterative proportional fitting allocates the industry totals across workers while preserving both margins. If p^b_{jf} is the QWI high-skill worker share for industry type f, the allocation satisfies

$$\sum_k L^b_{jfk} = L^b_{jf}, \qquad \sum_f L^b_{jfk} = L^b_{jk}, \qquad \frac{L^b_{jHH} L^b_{jLL}}{L^b_{jHL} L^b_{jLH}} = \frac{p^b_{jH}/(1-p^b_{jH})}{p^b_{jL}/(1-p^b_{jL})}. \tag{C.20}$$

The odds-ratio condition applies to positive flows; zero employment margins remain zero. Each fitted home–work flow is divided using its workplace's industry shares conditional on worker skill. The accessible outside location uses Census-2000 national industry shares conditional on education.

Let $\widetilde{w}^b_{jfk}$ be the county or CZ QWI earnings prior, and $\overline{w}^b_{jk}$ the existing broad-sector wage. Initial wages retain the QWI ratio between industry types and match the broad-sector payroll:

$$w^{b,0}_{jfk} = a^b_{jk} \widetilde{w}^b_{jfk}, \qquad a^b_{jk} = \frac{\overline{w}^b_{jk} L^b_{jk}}{\sum_{f \in \{H,L\}} L^b_{jfk} \widetilde{w}^b_{jfk}}. \tag{C.21}$$

This normalization applies to positive broad-sector employment. It preserves workplace payroll and aggregate earnings at each residence. Wages subsequently clear distinct workplace–industry–worker-skill labor markets. Skill cost shares are calibrated from each industry's payroll. Initial nontradable revenue follows its payroll and the common labor cost share. Establishment masses and baseline destination expenditure shares are divided according to these revenue shares. The two industry masses sum to the tract's initial establishment mass, and both industries enter one CES shopping basket. Each sector uses the common sector scale ν; its work and nonemployment choices use the separate common participation scale η.

C.5.2 Production cost shares

Data and industry aggregation. I use the 2012 detailed supply–use accounts of the U.S. Bureau of Economic Analysis (BEA). The source is the 2012 worksheet in `Use_SUT_Detail.xlsx`, together with the industry's classification in the workbook's `NAICS Codes` concordance.[37] All source amounts are in millions of current 2012 dollars, and industry gross output is valued at basic prices.

[37] The workbook is distributed in BEA's supply–use data package: https://apps.bea.gov/industry/release/zip/SUPPLY-USE.zip. The version retrieved on October 28, 2026, contains historical 2012 estimates organized using the workbook's 2017-NAICS concordance. The archived source and replication files preserve the data vintage.

Let $g \in \{T_H, T_L, N\}$ index the two tradable sectors and the nontradable composite, which pools N_H and N_L. The high-skill-intensive tradable group, T_H, contains NAICS 51, 52, 54, and 55. The low-skill-intensive group, T_L, contains manufacturing, NAICS 31–33. These labels describe broad production composites: both firm types employ both worker skills, and the grouping does not separately identify high-technology manufacturing. The nontradable composite contains retail trade, education, health care and social assistance, accommodation and food services, and other services (NAICS 44–45, 61, 62, 72, and 81). Real estate is excluded as an output sector but remains an input to the modeled industries. Retail output is measured by retail margins rather than the full value of merchandise sold.

Purchased intermediate inputs. For each industry i, let Y_i denote gross output, X_i total purchased intermediate inputs, and R_i purchases of real-estate services. These correspond to BEA rows `T018`, `T005`, and `531ORE`, respectively. For the set of industries $\mathcal{I}_g$ assigned to group g, define

$$Y_g = \sum_{i \in \mathcal{I}_g} Y_i, \qquad X_g = \sum_{i \in \mathcal{I}_g} X_i, \qquad R_g = \sum_{i \in \mathcal{I}_g} R_i. \tag{C.22}$$

The intermediate-input share is

$$\alpha_X^g = \frac{X_g - R_g}{Y_g}. \tag{C.23}$$

I sum monetary amounts before taking ratios, so the resulting share is weighted by industry gross output. Subtracting R_g keeps purchased real-estate services out of the outside-input composite and avoids assigning the same expenditure to both intermediate inputs and commercial space. The aggregation uses the BEA concordance, including the nonnumeric industry code `4B0000` for other retail activities.

Allocation between labor and commercial space. I allocate the remaining production-cost share according to

$$\alpha_L^g = 0.8(1 - \alpha_X^g), \qquad \alpha_C^g = 0.2(1 - \alpha_X^g). \tag{C.24}$$

Thus, $\alpha_L^g + \alpha_C^g + \alpha_X^g = 1$ and $\alpha_C^g/(\alpha_L^g + \alpha_C^g) = 0.2$ for each group. The common 80/20 split draws on the labor–floor-space calibration in Ahlfeldt et al. (2015), Hoelzlein (2023), and Miyauchi et al. (2025). Here, I apply that benchmark to the share remaining after purchased non-real-estate intermediates have been deducted. This retains sectoral differences in intermediate-input use without introducing additional group-specific labor–space parameters.

Purchased real-estate services alone do not measure the full economic cost of commercial premises, particularly when businesses own rather than rent their space. Conversely, gross operating surplus is not itself a purchase of intermediate inputs. These considerations motivate retaining the measured intermediate-input share while using a separate benchmark for the labor–space allocation.

The allocation is nevertheless a maintained approximation. The residual $Y_g - X_g + R_g$ contains

employee compensation, purchased real-estate services, gross operating surplus, and net production taxes, subject to published rounding. Operating surplus also includes depreciation and returns to capital other than commercial premises. Allocating this entire residual between labor and space therefore does not recover their separately observed national-account payments or imply that all operating surplus and taxes accrue to local workers and property owners. It maps the accounting residual into the model's restricted set of production factors.

Resulting shares. Table C.4 reports the source totals and implied allocations. I retain the unrounded ratios in the calibration.

Table C.4: 2012 BEA inputs and calibrated production shares

	T_H	T_L	N
Number of detailed industries	38	232	41
Gross output, Y_g	5,730,400	5,746,010	4,831,575
Purchased intermediates, X_g	2,376,329	3,855,036	1,943,393
Real-estate purchases, R_g	163,670	24,482	305,066
Non-real-estate inputs, $X_g - R_g$	2,212,659	3,830,554	1,638,327
Labor share, α_L^g	0.4911	0.2667	0.5287
Commercial-space share, α_C^g	0.1228	0.0667	0.1322
Intermediate-input share, α_X^g	0.3861	0.6666	0.3391
Space share within labor–space bundle	0.2000	0.2000	0.2000

Notes: Monetary amounts are millions of current 2012 dollars. T_H comprises NAICS 51, 52, 54, and 55; T_L comprises NAICS 31–33; and N comprises NAICS 44–45, 61, 62, 72, and 81. Intermediate-input shares are calculated from BEA accounts. The 80/20 allocation of the remaining share is imposed.

Interpretation in the model. The model treats the non-real-estate intermediate composite as supplied from outside the CZ at a fixed price. This is a sourcing assumption, not a geographic fact identified by the national BEA tables: observed purchases can include inputs supplied within the same CZ. Consequently, α_X^g measures purchased-input intensity in the data and expenditure leakage only under the model's sourcing closure.

For the tradable Cobb–Douglas technologies, the calibrated shares are production exponents. For the nontradable labor–space CES nest, the 80/20 allocation is a baseline cost-share target used to recover the distribution parameter at observed factor prices; expenditure shares can change when relative prices change. The allocation between the two worker skills within labor is calibrated separately.

With incumbent tradable output, commercial rents, and outside-input prices fixed, a proportional increase in both skill wages implies

$$\frac{\partial \log L_g}{\partial \log w} = -(1 - \alpha_L^g). \tag{C.25}$$

The corresponding elasticity magnitudes are approximately 0.509 for T_H and 0.733 for T_L. These are

implications of the calibrated technology and fixed-output closure, rather than elasticities estimated from the BEA accounts.

C.5.3 Establishment entry

I set $\epsilon_m = 0.40$, in line with the OLS coefficient of 0.39 from the county regression below. The regression is descriptive rather than causal, for the reasons given at the end of this subsection, but it shows that the maintained value matches how establishment counts have moved with service receipts across counties.

The free-entry condition in (20) gives the number of varieties in tract j, suppressing the sector index g,

$$m_j = \left(\frac{R_j^N}{\sigma_N \varphi_j c_j^N} \right)^{\epsilon_m}. \tag{C.26}$$

Aggregating to counties and taking logs and five-year differences,

$$\Delta \log m_{jt} = \epsilon_m \left(\Delta \log R_{jt}^N - \Delta \log c_{jt}^N \right) - \epsilon_m \Delta \log \varphi_{jt}. \tag{C.27}$$

I measure varieties by establishments, revenue by receipts, and unit-cost growth by wage growth, and I absorb common movements in input prices and entry costs with end-year effects δ_t:

$$\Delta \log m_{jt} = \epsilon_m \left(\Delta \log R_{jt}^N - \Delta \log w_{jt} \right) + \delta_t + e_{jt}. \tag{C.28}$$

The regression uses 9,798 five-year county–window observations, after aggregating the five service sectors and excluding NAICS 53, and gives a historical descriptive coefficient of 0.39 with a conventional OLS standard error of 0.01. The error contains $-\epsilon_m \Delta \log \varphi_{jt}$ and the gap between unit-cost and wage growth. Because entry also raises revenue by attracting spending, and receipts differ from the model's merchant-margin output, I read the estimate as descriptive rather than causal.

C.5.4 Accounting units and the property tax

Amounts are expressed in constant 2023 dollars. Individual income, rents, production, and tax revenue are measured over the same one-year accounting period. The baseline property tax rate is $t_a^0 = 0.28$. After entry, county rates adjust to finance service needs. Tax revenue purchases goods outside the modeled economy rather than providing cash rebates to residents. The event-specific local-benefit exercise is described in Appendix F.1.

I measure a sample-specific tax-to-housing-cost benchmark by adapting Couture et al. (2024, Section 4.1.3 and online Appendix C). They divide a population-weighted median of local median property-tax bills by a population-weighted median home value converted to an annual user cost. I apply this construction at the county level to the 44 distinct host counties of the 55 model events, counting each county once. The 2014 ACS five-year release covers 2010–2014; its end year equals the median event year, so this is a common-period benchmark rather than a uniformly pre-entry

measure. County population, median annual real-estate taxes, and median owner-occupied home value come from ACS tables B01003, B25103, and B25077, respectively. All sample counties have valid estimates. Following the authors' conversion for 2014, I annualize home values at 4.6%:

$$\widehat{t}^{\,\mathrm{ACS}} = \frac{\mathrm{wmed}_{c,\,\mathrm{pop}_c}(\text{median annual tax}_c)}{0.046\ \mathrm{wmed}_{c,\,\mathrm{pop}_c}(\text{median home value}_c)} = \frac{\$2{,}432}{\$7{,}700.40} = 31.6\%.$$

The underlying population-weighted median county home value is $167,400. The numerator and denominator are aggregated separately; this is not a median of household tax-to-cost ratios. The common dollar deflator cancels from the ratio. Using aggregate taxes and aggregate home values instead (ACS B25090 and B25082), with the same annualization, gives 26.1%. For comparison, taxes are 17.4% of observed annual owner costs (twelve times ACS B25089); those costs include mortgage payments, taxes, utilities, insurance, and relevant fees, so this is a different denominator. The ACS measures owner-occupied housing and taxes paid to all taxing jurisdictions, not county-government receipts alone or commercial tax rates. These measurements are empirical benchmarks; the equilibrium calculations use $t_a^0 = 0.28$ on pretax rent as the baseline rate, with counterfactual rates determined by the county budget.

C.6 Building the pre-entry economy

This subsection supports Section 5.3.

C.6.1 Baseline allocations and preference parameters

The baseline allocation covers individuals in private tradable and service employment and in nonemployment. The observed residence and workplace margins are balanced using the commuting gravity weights. Inbound and outbound flows connect the CZ to the neighboring region. LODES does not report commuting flows by education, so the boundary flows by skill use LODES earnings bands as a proxy: the top band, with monthly earnings above $3,333, stands in for high-skill workers and the two lower bands for low-skill workers. For each tract and skill, the earnings-band flows set the share of jobs held by residents of the CZ rather than inbound commuters and the share of residents who work outside the CZ. These shares are applied to the QWI workplace totals by skill, which are measured by education. The additional outside-residence/outside-workplace worker stock is neighboring-CZ employment net of existing inbound workers. Nonworker counts and national population stocks complete the finite-population account.

Let $n_{\ell k}^0$ be the baseline mass of skill k in choice ℓ, including the national choice. For a nest v whose members $\mathcal{C}(v)$ are location pairs or broader choices, write M_{vk}^0 for the baseline mass of all choices in the nest. Let $\mathcal{I}(\ell)$ collect the conditional choices specifying ℓ. The nest's baseline conditional probabilities are

$$\pi_{u|v,k}^0 = \frac{M_{uk}^0}{M_{vk}^0}, \qquad M_{vk}^0 = \sum_{u \in \mathcal{C}(v)} M_{uk}^0, \qquad n_{\ell k}^0 = \mathcal{N}_k \prod_{(u,v) \in \mathcal{I}(\ell)} \pi_{u|v,k}^0. \tag{C.29}$$

On choices with positive baseline mass, these conditional probabilities reproduce participation, sector, area, and commuting allocations. Location-pair commuting probabilities within the CZ are predicted by the fitted gravity system; they do not reproduce every realized bilateral count. The amenities recover these fitted probabilities in levels. At broader nests, (B.32) recovers the preference parameters given the baseline inclusive values. The shape parameters govern responses to relative utility. The scale order and the placement of the inbound commuting nest are in Appendix B.2.

Location amenities. Let a^0_{ik} and b^0_{jk} be the fitted residence and workplace factors for skill k, and ϑ^0_{jgk} the share of workplace employment in sector g. Within the CZ, fitted commuting flows are allocated across industries using these workplace shares:

$$F^0_{ijgk} = a^0_{ik} b^0_{jk} \vartheta^0_{jgk} \exp(-\chi \kappa t_{ij}), \qquad i, j \in \mathcal{J}. \tag{C.30}$$

The workplace shares sum to one across employment sectors. Because industry allocation depends on workplace and skill, the residential factor remains common across sectors. Let $\gamma = (1 - \alpha_H) s_N$ and let $y^0_{jgk} = w^{g,0}_{jk} + D^0_k$ be disposable income for a CZ resident employed at j in sector g. One normalization of the location amenities is

$$\begin{aligned} B^R_{ik} &= (a^0_{ik})^{1/\chi} (q^{R,0}_i)^{\alpha_H} (P^0_i)^{\rho\gamma}, \\ B^W_{jgk} &= (b^0_{jk} \vartheta^0_{jgk})^{1/\chi} \frac{(P^0_j)^{(1-\rho)\gamma}}{y^0_{jgk}}. \end{aligned} \tag{C.31}$$

Substitution gives $(\Phi^0_\ell)^\chi = F^0_{ijgk}$ on internal working choices. Hence normalized utility powers reproduce the fitted location-pair probabilities without a further bilateral preference parameter. The broader preference parameters are expressed in the corresponding value units using (B.32).

Boundary flows retain their separately measured totals. For an existing working choice involving the outside location, define its fixed amenity composite by

$$\mathcal{B}_\ell = (F^0_{ijgk})^{1/\chi} \frac{(q^{R,0}_i)^{\alpha_\ell} (\mathcal{P}^0_\ell)^{\gamma_\ell}}{y^0_\ell}, \qquad i = o \text{ or } j = o. \tag{C.32}$$

These are the outbound, inbound, and outside–outside composites in (B.1). They index only the local origin or local destination, or the single outside worker stock. They introduce no unrestricted residual for an internal residence–workplace pair. For nonemployment, choose B^U_{ik} so that $(\Phi^0_\ell)^\chi = U^0_{ik}$ for each nonworking choice at i; conditional residential probabilities then equal $U^0_{ik} / \sum_{a \in \mathcal{I}_r} U^0_{ak}$ in every sector. Broader nest probabilities allocate the nonworkers across sectors as in (B.35).

These inversions apply to choices with positive baseline mass.

C.6.2 Commuting gravity

Consider internal commuters in one skill–sector group and suppress its indices. The residence and workplace amenities are the fitted factors recovered in (C.31); the workplace amenity includes

the sector dimension suppressed here. Let $T = D_{ck}$ be the common local dividend and set $\gamma_c = (1 - \alpha_{H,c})s_N$. Cobb–Douglas maximization gives

$$\Phi_{ij} = \underbrace{\frac{B_i^R}{(q_i^R)^{\alpha_{H,c}} P_i^{\rho\gamma_c}}}_{\mathcal{R}_i} \underbrace{\frac{B_j^W (w_j + T)}{P_j^{(1-\rho)\gamma_c}}}_{\mathcal{W}_j} e^{-\kappa t_{ij}}. \tag{C.33}$$

The normalized-power probabilities within the bilateral choice group are proportional to Φ_{ij}^{χ}. Thus its internal flows have the gravity representation

$$F_{ij} = a_i b_j e^{-\chi\kappa t_{ij}}, \qquad \log F_{ij} = \delta_i + \zeta_j - \chi\kappa t_{ij}. \tag{C.34}$$

Writing M for the group's commuter mass and $Z = \sum_{a,b} (\mathcal{R}_a \mathcal{W}_b)^{\chi} \exp(-\chi\kappa t_{ab})$, the factors can be chosen as $a_i = M\mathcal{R}_i^{\chi}/Z$ and $b_j = \mathcal{W}_j^{\chi}$. At baseline, positive fitted factors solve the margin equations

$$O_i^0 = a_i^0 \sum_j b_j^0 e^{-\chi\kappa t_{ij}}, \qquad D_j^0 = b_j^0 \sum_i a_i^0 e^{-\chi\kappa t_{ij}}, \tag{C.35}$$

where O_i^0 and D_j^0 are the calibrated resident-worker and workplace margins for this group. One normalization fixes their otherwise arbitrary relative scale. Given baseline budgets and prices, the fitted factors identify the corresponding residence and workplace amenity composites. The fitted gravity system reproduces these tract margins and predicts the bilateral counts. An observed zero on a feasible link is compatible with a positive fitted expected flow.

To see how the model solves a counterfactual, let π_{ij}^0 be the normalized baseline fitted flows. Within this group,

$$\pi_{ij}^1 = \frac{\pi_{ij}^0 \exp(\chi \Delta v_{ij})}{\sum_{a,b} \pi_{ab}^0 \exp(\chi \Delta v_{ab})}, \tag{C.36}$$

$$\Delta v_{ij} = \log \frac{w_j^1 + T^1}{w_j^0 + T^0} - \alpha_{H,c} \log \frac{q_i^{R,1}}{q_i^{R,0}} - \gamma_c \left[\rho \log \frac{P_i^1}{P_i^0} + (1 - \rho) \log \frac{P_j^1}{P_j^0} \right]. \tag{C.37}$$

The common commute cost and fixed amenities cancel in the value change; they remain encoded in π_{ij}^0. Market clearing determines wages, rents, shopping prices, and dividends, while the higher nests determine the group's counterfactual mass. This update uses fitted probabilities, so a zero realized survey or administrative count need not prevent a positive model flow after entry.

Residence and workplace factors incorporate income, prices, amenities, and group normalization. Higher choice nests determine the group's mass; boundary-flow calibration connects it to the outside economy. Baseline internal commuting probabilities are fitted to travel times and tract margins. Bilateral counts are predictions, with positive expected flows on feasible links between positive margins. This covariate-based approach follows the calibration principle examined by Dingel and

Tintelnot (2026), who distinguish estimated probabilities from noisy realized counts in granular commuting data. Exact-hat updates can use these fitted probabilities.

The functional-form restriction can be seen by substituting the auxiliary Stone–Geary demand into the same joint-choice structure. Resources in (C.33) would become $w_j + T - q_i^R \bar{h}_c$, and log flows would contain the additional term $\chi \log(w_j + T - q_i^R \bar{h}_c)$. This origin–destination interaction generally cannot be absorbed by separate fixed effects. It arises from purchasing power: the same minimum housing bill absorbs a larger fraction of resources at a lower-paying workplace. The Cobb–Douglas specification makes that purchasing-power relationship multiplicatively separable while retaining individual and ownership budget consistency.

C.6.3 Production costs and payroll

Appendix B.7 derives unit cost and factor demands. The following calibration applies to each service sector $g \in \mathcal{G}_N$.

The baseline shares $(\alpha_L^N, \alpha_C^N, \alpha_X^N) = (0.53, 0.13, 0.34)$ are the nontradable-composite shares in Table C.4, common to N_H and N_L. They are the shares α_L^g, α_C^g, and α_X^g of both service sectors $g \in \mathcal{G}_N$ in (17). Because production is Cobb–Douglas, these shares hold in every tract and at any factor prices, so no tract-level inversion is needed. The outside category is purchased intermediate inputs other than real-estate services. Property taxes enter through the tax-inclusive commercial rent.

The model's labor bill equals payroll:

$$W_{jg} = \sum_k w_{jk}^g L_{jk}^g = \alpha_L^g C_{jg} = \alpha_L^g R_{jg}. \tag{C.38}$$

The same wage determines firms' labor costs and individuals' employment income. The labor share above is applied directly to payroll.

Observed revenue, unit cost, and establishment mass determine the baseline fixed-input shifter through the zero-profit condition:

$$\varphi_{jg} = \frac{R_{jg}^0}{\sigma_N c_{jg}^0 (m_{jg}^0)^{1/\epsilon_m}}. \tag{C.39}$$

The inversion uses the specified shopping elasticity, production technology, and property tax. Reported entrant wages provide a validation comparison for the common workplace–skill tradable wages used in (24).

C.6.4 Service revenue and productivity inversion

The inversion recovers the service productivity and operating-cost shifters that make household spending, the shopping network, and service revenue consistent in the baseline equilibrium. It adapts the approach of Oh and Seo (2023) and proceeds in four steps. Revenue targets are inferred

from payroll. The service expenditure share is set so that household spending covers these targets in aggregate. Destination productivities allocate that spending across tracts. Operating-cost shifters then reconcile revenue per establishment with free entry. The inversion is performed separately for each event's commuting zone, and I suppress the CZ index below. Superscript 0 denotes the baseline, i indexes shopping origins, j indexes destinations, and $g \in \mathcal{G}_N = \{N_H, N_L\}$ indexes the two service sectors.

Revenue targets. Service sales are not observed by tract, so I infer them from payroll. Let W_{jg}^0 denote annual service payroll and m_{jg}^0 the number of establishments in tract j and sector g. Variable production and establishment operation use the same technology, so under free entry labor payments equal the labor share α_L^N of revenue (equation (C.38)). The revenue targets are therefore

$$R_{jg}^0 = \frac{W_{jg}^0}{\alpha_L^N}, \qquad R_j^0 = \sum_{g \in \mathcal{G}_N} R_{jg}^0, \qquad R^0 = \sum_j R_j^0. \tag{C.40}$$

These targets follow from payroll and the production specification; they are not independently observed sales.

Aggregate expenditure. Let Q_i^0 denote the income remaining after housing, including dividends, that is assigned to shopping origin i, before the service expenditure share is applied. An employed individual assigns the fraction ρ of this income to her residence and $1-\rho$ to her workplace, and a nonworker assigns all of it to her residence. Inward commuters therefore contribute to the budgets of their workplace origins, while income assigned to origins outside the CZ is excluded from Q_i^0. Service expenditure originating at i is $\mathcal{E}_i^0 = s_N Q_i^0$.

Let $\bar{\pi}^O$ be the measured share of these budgets spent outside the CZ, weighted by spending. The service expenditure share at which household spending covers aggregate revenue is

$$s_N^{\text{req}} = \frac{R^0}{(1-\bar{\pi}^O)\sum_i Q_i^0}. \tag{C.41}$$

Because individuals cannot spend more than their income after housing on services, I cap the share at one and define the pooled external budget $\mathcal{E}^{\text{ext}}$ as the remaining shortfall,

$$s_N = \min\{s_N^{\text{req}}, 1\}, \qquad \mathcal{E}^{\text{ext}} = R^0 - (1-\bar{\pi}^O) s_N \sum_i Q_i^0. \tag{C.42}$$

Thus $\mathcal{E}^{\text{ext}} = 0$ whenever household spending can cover the revenue targets. Otherwise, $\mathcal{E}^{\text{ext}}$ is the smallest additional budget that reconciles the aggregate accounts. It represents purchases financed outside the modeled household budgets.

Spatial allocation and productivity. Given total spending, the inversion recovers where it goes. Let Γ_{ij} and Γ_i^O denote the initial shopping weights for destination j and for the outside

alternative. These weights carry the shopping network's travel-cost structure and its baseline normalization of destinations. For a vector of destination multipliers $\delta = (\delta_j)_j$, the shares of origin i's spending that go to j and outside the CZ are

$$\pi_{ij}(\delta) = \frac{\Gamma_{ij}\delta_j}{\Gamma_i^O + \sum_{j'} \Gamma_{ij'}\delta_{j'}}, \qquad \pi_i^O(\delta) = \frac{\Gamma_i^O}{\Gamma_i^O + \sum_{j'} \Gamma_{ij'}\delta_{j'}}. \tag{C.43}$$

The outside weight is fixed, which pins down the level of the destination multipliers.

The external budget follows the distribution of household purchases, so households finance the fraction $1 - \mathcal{E}^{\text{ext}}/R^0$ of each tract's revenue. I recover δ by solving

$$\sum_i \mathcal{E}_i^0 \pi_{ij}(\delta) = \left(1 - \frac{\mathcal{E}^{\text{ext}}}{R^0}\right) R_j^0 \qquad \text{for each active destination } j. \tag{C.44}$$

Summing these equations across destinations shows that

$$\frac{\sum_i \mathcal{E}_i^0 \pi_i^O(\delta)}{\sum_i \mathcal{E}_i^0} = \bar{\pi}^O, \tag{C.45}$$

so the inversion preserves aggregate shopping leakage while allowing each origin's outside share and its bilateral destination shares to change.

I solve the system by iterative proportional fitting. Each iteration scales every destination multiplier by the ratio of its target to its predicted household revenue, with a common normalization that keeps the outside weight fixed. The calculation keeps the shopping network's existing support, so every destination with a positive revenue target must be reachable from origins with positive budgets.

The multipliers map into productivity through CES demand. With elasticity $\sigma_N > 1$, a destination's attractiveness is proportional to its price raised to the power $1 - \sigma_N$, and at fixed input prices and establishment counts, higher productivity lowers price proportionally. Relative to the initial normalization, the recovered multiplier therefore corresponds to

$$\frac{A_{jg}^0}{A_{jg}^{\text{ref}}} = \delta_j^{1/(\sigma_N - 1)}, \qquad \frac{p_{jg}^0}{p_{jg}^{\text{ref}}} = \delta_j^{-1/(\sigma_N - 1)}. \tag{C.46}$$

A destination that receives more revenue than the shopping network alone would send it is thus inferred to have more productive service firms, whose lower prices draw more spending. The productivity adjustment is common to the two service industries within a tract, and their baseline shares of the tract's spending are set separately as

$$\vartheta_{jg}^0 = \frac{R_{jg}^0}{R_j^0}. \tag{C.47}$$

Together, the tract multipliers and industry spending weights reproduce each tract–industry revenue target. The recovered productivities are conditional on the maintained shopping preferences

and cost specification, and they do not separate physical productivity from unobserved attractiveness of a destination.

Operating costs and baseline consistency. With CES markup pricing, variable profits equal revenue divided by σ_N, so free entry determines baseline operating expenditure per establishment,

$$f_{jg}^0 = \frac{R_{jg}^0}{\sigma_N m_{jg}^0}. \tag{C.48}$$

I recover the operating-cost shifter that delivers this expenditure at baseline input prices and productivity (equation (C.39)), which preserves the observed establishment counts. Baseline amenity normalizations absorb the resulting change in service price levels, so the inversion also preserves baseline worker allocations. The baseline equilibrium therefore reproduces the employment, payroll, revenue, and establishment-count targets.

Counterfactual expenditure and revenue. The recovered productivity and operating-cost shifters, the service expenditure share, and the total external budget are held fixed in counterfactuals, while establishment counts and prices are endogenous. Let hats denote ratios to baseline values, and let π_{ij}^0 and $\pi_i^{O,0}$ be the fitted baseline shopping shares. The counterfactual share of origin i's spending received by tract j and sector g is

$$\pi'_{ijg} = \frac{\pi_{ij}^0 \vartheta_{jg}^0 \widehat{m}_{jg} \widehat{p}_{jg}^{1-\sigma_N}}{\pi_i^{O,0} + \sum_{j',g'} \pi_{ij'}^0 \vartheta_{j'g'}^0 \widehat{m}_{j'g'} \widehat{p}_{j'g'}^{1-\sigma_N}}. \tag{C.49}$$

Outside opportunities stay fixed, but the share spent outside the CZ changes as inside prices and varieties change.

Given counterfactual origin budgets $\mathcal{E}'_i$, household purchases and total revenue are

$$\begin{aligned} R_{jg}^{HH\prime} &= \sum_i \mathcal{E}'_i \pi'_{ijg}, \\ R'_{jg} &= R_{jg}^{HH\prime} + \mathcal{E}^{\text{ext}} \frac{R_{jg}^{HH\prime}}{\sum_{j',g'} R_{j'g'}^{HH\prime}}. \end{aligned} \tag{C.50}$$

The external budget is fixed in nominal terms, but its allocation follows current household purchases. The revenue targets are imposed only in the baseline. In counterfactuals, revenues are determined by equilibrium spending.

C.7 Equilibrium computation and verification

The current solver uses sector–region ordering for high-skill workers and region–sector ordering for low-skill workers, followed by the participation and workplace nests, with tract-specific residential supply.

Individual choices are evaluated using fixed baseline predicted probabilities and changes in utility, as in (B.33). Fixed amenities and commuting costs cancel from utility ratios and remain encoded in those probabilities. Equation (B.46) normalizes the available choices in each equilibrium.

Each equilibrium jointly solves for wages by sector and skill, residential and commercial rents, service establishment mass, and portfolio dividends, all represented in logs. A matrix-free Newton method uses GMRES steps and a line search. Continuation in the size of the entry is available when a direct full-entry solve fails. Reported equilibria satisfy the full counterfactual specification.

Saved reference states are re-verified before reuse, and newly solved states satisfy equilibrium and accounting residual tolerances of 10^{-8}. Each saved state allocates positive disposable income to every available individual choice.

Dollar equivalent variation is evaluated after the equilibrium solve. A bracketed scalar root-finder finds the transfer at reference prices that matches the entry opportunity index; this calculation does not solve the equilibrium again. Safeguarded interpolation and bisection give the same equivalent-variation solution. Initializing nearby counterfactual exercises from verified states reduces the equilibrium iterations without relaxing the residual checks.

No recorded cash award enters the current equilibria. The new-property exemption and endogenous county property-tax rates do enter the equilibrium. Shopping rules and the combined home-shopping/inbound counterfactual use their own no-entry equilibria.

Of 226 public events, 216 have certified observed-entry results. Nine remain blocked by source data or metadata and one lacks a certified baseline. The reported observed-entry sample therefore spans 95 CZs.

Numerical convergence establishes that the reported states satisfy the specified equations; it does not establish uniqueness or empirical validity.

C.8 Exposure for the validation scorecard

Updated results pending. A current-model scorecard matching the empirical worker-source and exposure definitions has not yet been produced. The available baseline-residence, neighboring-tract proxy is a preliminary diagnostic and is not substituted for that validation.

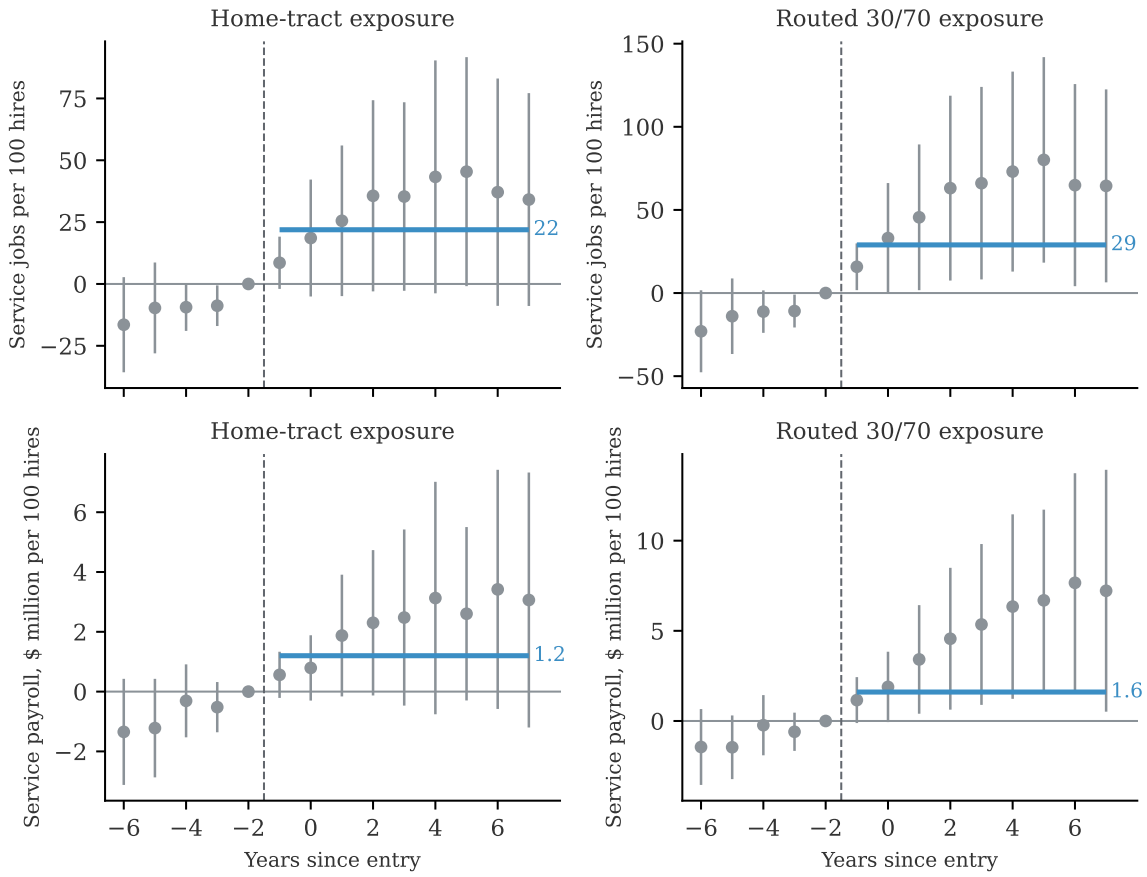
D Validation: Supporting Comparisons

This appendix reports the two comparisons that Section 5.4 cites but does not show, the Fact 3 event study for all nontradables and the stacked contrast between the top and bottom quartiles of exposed tracts.

All nontradables. Figure D.1 repeats Figure 7 with the broader service definition of the release, which adds education and health (NAICS 61 and 62) to retail, food and accommodation, and other services. The model's response is 22 jobs per 100 exposed hires under the home-tract exposure against a years-three-to-seven mean of 39 in the data, and 29 against 70 under the hyperlocal

exposure; payroll is $1.2 million against $2.9 million and $1.6 million against $6.7 million. The ratios are 0.56 and 0.42 for jobs and 0.41 and 0.24 for payroll, against 0.91 and 0.68 for core jobs and 0.51 and 0.39 for core payroll on the same years-three-to-seven means in the main text. The model lies inside the annual 95 percent interval at every horizon under the home-tract exposure and at every horizon but year five of the payroll design under the hyperlocal exposure. The gap that opens when education and health are added is the one the main text attributes to how those sectors are funded and sited; the model treats them as demand-driven services and has no separate margin for them.

Figure D.1: All nontradables where the entrant's workers live: empirical event study and the model



Notes: As Figure 7, for all nontradables (NAICS 44 to 45, 61, 62, 72, 81). Circles are the released continuous event-study coefficients per 100 exposed hires with 95 percent confidence intervals; year −2 is the reference and hiring begins in year −1 (dashed line). The blue line is the model's static response from the same design on all 216 events.

The stacked quartile contrast. The release also reports a stacked design that compares the change in service employment between the top and bottom quartiles of exposed tracts within each event. Table D.1 places the model's analog beside it. The model's contrast is a tenth of the data's under either exposure. The ratio decomposes exactly into the ratio of exposure gaps, how

much more exposed the top quartile is than the bottom in the model relative to the data, and the ratio of responses per unit of gap. For the hyperlocal exposure at year seven the model-to-data ratio of 0.11 is 0.38 times 0.30; for the home-tract exposure at year three it is 0.10, equal to 0.29 times 0.34. Both factors fall short. The exposure gap is smaller in the model because the data concentrate hires' residences in a few tracts near each site while the model spreads them more evenly across the commuting zone, and the response per unit of gap is smaller for the same reason the continuous payroll response is, the model's service jobs per dollar of local spending. The stacked design therefore tests the concentration of residences as much as the demand mechanism, and the continuous design in the main text is the one that isolates the latter.

Table D.1: Stacked top-minus-bottom quartile contrast: released estimates and the model

Exposure	Outcome	Data, mean of years 3–7	Model	Ratio
Home tract	Service jobs per 100 hires	39.2	3.6	0.09
Home tract	Service payroll ($ million)	3.39	0.20	0.06
Hyperlocal 30/70	Service jobs per 100 hires	30.5	4.0	0.13
Hyperlocal 30/70	Service payroll ($ million)	3.19	0.22	0.07

Notes: Released stacked estimates for all nontradables from files 1A and 1C (30 and 40 commuting zones), averaged over years three to seven; the release gives no cross-year covariance, so no interval is reported for the mean. Model values are equal-event averages of the top-minus-bottom quartile contrast on the 53 single-event commuting zones, for the model's full service sector, so they are broad-outcome analogs rather than core-sector tests. The decomposition in the text uses the released and model exposure gaps at the stated horizon.

E Employment, Welfare, and Incidence: Supporting Results

This appendix supports Section 6 in the order of its subsections. The observed sample contains 216 events in 95 CZs. Observed-event summaries weight events equally. All money is annual 2023 dollars; "resident" refers to the model's adult population, including nonworkers aged 25–64; other-age nonworkers are excluded. Appendix F reports the award-accounting and location comparisons for Section 7.

E.1 Employment and the payroll accounting chain

This subsection supports the employment and payroll results in Section 6. Let J_e be the promised job target and D_e the realized employment increase in the entrant's tract and industry without a property-tax exemption, $\Delta T_e^{\text{other}}$ the change in other tradable jobs, and ΔN_e the change in service jobs. The observed-site main model satisfies

$$\Delta L_e^{\text{net}} = D_e + \Delta T_e^{\text{other}} + \Delta N_e, \qquad M_e^N = 100\frac{\Delta N_e}{J_e}, \qquad M_e^{\text{net}} = 100\frac{\Delta L_e^{\text{net}}}{J_e}. \tag{E.1}$$

The service multiplier per net tradable job is $100\Delta N_e/(D_e + \Delta T_e^{\text{other}})$ when its denominator is positive. Productivity is calibrated to J_e; numerical solution tolerances allow a small difference

between D_e and J_e. The reported median is taken over the event-level ratio, rather than a ratio of separately reported medians.

Let ΔP_e be the additional payroll at the entrant's site. Each ledger entry is $100\Delta X_e/\Delta P_e$. Resident earnings equal workplace payroll less inbound workers' earnings plus residents' earnings outside the CZ. Disposable income adds nonemployment payments and private rental dividends. Resident service budgets are $s_{N,c}(1-\alpha_i)$ times income at each residence. Home- and workplace-origin budgets can differ from resident budgets because workers cross the CZ border. Allocating these origin budgets over shopping destinations gives spending at service firms in the CZ. Inside-origin shopping has no outside service destination in this specification; outside residents still spend their home-origin budgets outside the CZ. Identities are checked in levels within events; medians across events are not an additive accounting chain.

Table E.1 reports the full chain per \$100 of additional payroll in the shocked tract and industry. Residential rental dividends add about \$10, and small changes in nonwork payments bring the increase in resident disposable income to about \$53. Housing expenditure rises by about \$11, and spending on tradable goods absorbs part of the remainder, leaving residents' service budget about \$26 higher. Net imports of workplace-origin shopping budgets add about \$5. With no outside service destination for inside origins, none of these budgets leaks through that shopping margin. Spending at service firms inside the CZ rises by about \$31. These figures are medians of event-level ratios, so the reported amounts need not add together; the accounting identities hold within each event.

Table E.1: Payroll, resident income, and service spending per \$100 of additional payroll at the entrant's site

Account	Median	P10	P90
CZ workplace payroll	81.24	71.44	92.00
Resident labor earnings	40.84	29.90	53.27
Nonemployment payments	0.73	0.36	1.28
Rental dividends	10.36	7.75	14.05
Resident disposable income	52.63	39.03	69.09
Housing expenditure	10.84	7.63	14.88
Resident service budget	25.66	16.58	36.77
Net workplace budget imports	5.45	3.47	7.62
Local-origin service budget	30.88	22.90	42.14
Spending outside	0.00	0.00	0.00
Spending at local service firms	30.88	22.90	42.14

E.2 Welfare measure

This subsection defines the welfare measure used throughout Section 6. For skill k and place of residence a, a county or the CZ, let V_{ak}^t be the model's inclusive value over feasible residential and

employment opportunities. The cash equivalent EV_{ak} solves

$$V^0_{ak}(y^0 + \mathrm{EV}_{ak}, p^0) = V^1_{ak}(y^1, p^1), \qquad \mathrm{EV}_a = \sum_k N^0_{ak}\mathrm{EV}_{ak}, \qquad \overline{\mathrm{EV}}_a = \frac{\mathrm{EV}_a}{\sum_k N^0_{ak}}. \tag{E.2}$$

The transfer is uniform across the baseline alternatives of a given skill. Prices and primitives on the left are held at their baseline values. Reported welfare per entrant job is EV_c/J_e; its wage equivalent divides that ratio by the employment-weighted baseline wage in the entrant's tract and industry. Both ratios are formed within events before reporting their medians. The measure values a place's opportunities and does not track the welfare of a fixed cohort. Regional indices are nonlinear, so host-county and rest-of-CZ welfare are evaluated directly, not obtained by subtracting or adding tract welfare totals.

E.3 Commuting and shopping restrictions

Figure 8 compares four restrictions with the main model. The home-origin specification sends all shopping trips from residences while retaining the shopping destination network. The home-tract-or-nearest specification instead sends each resident's service spending to the home tract if it has service supply, or to the nearest supported service tract by travel time. Both shopping specifications have their own solved no-entry equilibrium; the productivity and preference primitives are not re-inverted to fit the restricted baseline.

The no-migration specification holds the inside-CZ and adjacent-region populations at their no-entry values by worker skill. It allows residential reallocation within each region, participation, sector switching, and commuting. The migrant-only benchmark closes inbound commuting in both equilibria and fixes native participation. A new migrant cohort fills the added target jobs and chooses residences inside the CZ, while native employment in the target location-sector is held at its no-entry level. Other native job and residential choices remain endogenous.

Each comparison uses the same labor parameters as the main model, has no property-tax exemption, and excludes outside service destinations for inside shopping origins. The first three restrictions recalibrate productivity to raise employment in the target location-sector by the promised jobs relative to their no-entry equilibrium. The migrant-only benchmark retains the main-model productivity shock and instead fixes the size of the committed migrant cohort to achieve the same job target. Thus all four comparisons hold added target jobs fixed. Table E.2 reports quantiles of event-level differences in service jobs per 100 promised jobs. Figure 8 instead reports percentage differences between the restricted and main-model medians on the same event sample; these are not medians of individual percentage changes.

E.4 Employment adjustment under the restrictions

Table E.3 reports service employment, losses at other tradable producers, and net workplace employment. Table E.4 compares each restriction with the main model within the same events. Losses

Table E.2: Paired service-employment effects relative to the main model

Paired service-job effect	Median	P10	P90
Home origins only / County	-4.83	-7.41	-2.53
Home origins only / CZ	-4.82	-7.77	-2.97
Home tract or nearest / County	-4.27	-10.55	2.32
Home tract or nearest / CZ	-6.46	-10.59	-3.61
No migration / County	-3.30	-9.84	-0.59
No migration / CZ	-11.21	-15.48	-6.67
Migrant-only benchmark / County	4.23	-0.32	12.57
Migrant-only benchmark / CZ	16.26	10.88	23.86

Notes: Each row subtracts the same event's main-model service response, in jobs per 100 promised jobs. County denotes the host county. The first three restrictions have 216 paired events; the migrant-only benchmark has 206. P10 and P90 are the 10th and 90th percentiles of the paired differences.

at other tradable producers exclude the expanded location-sector. They enter net jobs with the opposite sign. Regional-population and native employment restrictions are quantity constraints, not re-estimated elasticities. All entries are static equilibrium changes relative to the appropriate no-entry comparator.

Table E.3: Mechanism restrictions: median equilibrium responses

Specification	Events	County service	CZ service	Other tradable loss	CZ net jobs
Main model	216	15.61	29.10	49.17	80.00
Home origins only	216	9.64	23.56	46.13	77.64
Home tract or nearest	216	10.91	21.77	44.32	76.93
No migration	216	11.53	17.26	58.33	59.16
Migrant-only benchmark	206	19.27	44.38	21.36	122.64

Notes: All job responses are per 100 promised jobs. Other tradable loss is positive for contraction elsewhere in the CZ. The main-model row covers 216 events. The migrant-only benchmark covers 206 certified events; ten unfinished or unsuccessful cases are omitted, not assigned zero effects. The other restrictions cover all 216 events. The following table uses an identical event sample for each paired comparison.

Table E.4: Mechanism restrictions: median paired changes from the main model

Specification	Events	County service	CZ service	Other tradable loss	CZ net jobs
Home origins only	216	-4.83	-4.82	-2.37	-2.36
Home tract or nearest	216	-4.27	-6.46	-3.67	-2.80
No migration	216	-3.30	-11.21	9.09	-20.29
Migrant-only benchmark	206	4.23	16.26	-28.00	45.12

Notes: Each restriction's entry-minus-no-entry response is compared with the same event's main-model response. Entries are medians of paired differences, not differences of medians. Units are jobs per 100 promised jobs. A positive change in other tradable loss means greater displacement. These restrictions are separate experiments, not additive components.

E.5 Spatial incidence, exposures, and regression checks

Table E.5 uses baseline commuting and shopping probabilities to measure exposure to the expansion's payroll. For event e, let j_e and g_e denote the expanded workplace and industry, $\mathcal{I}_e$ the tracts in its CZ, and O the outside residence. Let L^0_{e,ij_eg_ek} be baseline workers of skill k who live in i and work in the target location-sector, and let w^0_{e,j_eg_ek} be their cash wage. The share of the target location-sector's payroll received by residents of i is

$$p_{ei} = \frac{\sum_k w^0_{e,j_eg_ek} L^0_{e,ij_eg_ek}}{\sum_{r\in\mathcal{I}_e\cup\{O\}} \sum_k w^0_{e,j_eg_ek} L^0_{e,rj_eg_ek}}. \tag{E.3}$$

These are fitted baseline worker flows, not the expansion's realized hires. Including outside residents in the denominator preserves payroll leakage through inbound commuting.

Write $b_{ei} = s_{N,c(e)}(1-\alpha^0_{H,ei})$ for the fraction of income allocated to services after housing at residence i, using its baseline housing expenditure share $\alpha^0_{H,ei}$. Let $S^0_{e,ij}$ be the baseline probability that shopping from origin i reaches tract j, after the service-productivity inversion. With home-origin share ρ, the two exposures are

$$H_{ej} = 100\rho \sum_{i\in\mathcal{I}_e} p_{ei} b_{ei} S^0_{e,ij}, \tag{E.4}$$

$$W_{ej} = 100(1-\rho) S^0_{e,j_ej} \sum_{i\in\mathcal{I}_e\cup\{O\}} p_{ei} b_{ei}. \tag{E.5}$$

Outside residents' home-origin purchases remain outside the CZ; their work-origin purchases enter through W_{ej}. Both measures are dollars reaching a destination per \$100 of gross baseline payroll in the target location-sector. They do not measure the actual increase in residents' income or service spending. The shopping probabilities reproduce baseline household-financed service revenue; any pooled external service demand is excluded when recovering that household revenue target.

The regressions use the same complete-case sample in every column, excluding the expansion tract and tracts without a positive measurable baseline service wage. Controls are log one plus service employment, log average service cash wages, and log residential rents. Average wages combine both service industries and both worker skills. Event-by-county fixed effects remove the county mean for each event; observations receive equal weight. Standard errors use CZ-clustered CR1 covariance, accounting for absorbed effects in the finite-sample correction, and two-sided t_{94} tests.

As a prediction check, I remove event-by-county means and estimate the slopes on all but one CZ, repeating for each CZ. Applying these slopes to the omitted CZ gives a pooled within-county R^2 of 77.20 percent for the full specification, compared with 77.63 percent in sample. This checks whether the relationship carries across simulated CZs; it does not constitute validation against independent observed outcomes.

Figures 9 and 10 illustrate the current model's spatial incidence using the same no-exemption specification. The maps distinguish changes at workplaces from changes in residents' income.

To explain why service jobs grow in certain tracts, I regress each tract's service-job gain on its baseline service employment, wages, and rents, and its exposure to the entrant's payroll either through the home or work origin shopping (Table E.5). Both exposures come from the baseline equilibrium, before the productivity shock changes the distribution of individuals, wages and prices. I assign the baseline payroll of the entrant's location and sector to the tracts where its workers live using baseline commuting flows, keep the share residents spend on services after housing, and send that spending to destination tracts using baseline shopping probabilities, starting from workers' homes for home-route exposure and from the entrant's tract for work-route exposure. I find that pre-period exposure explains most of this variation. Comparing tracts within the same event and county, and leaving out the entrant's own tract, baseline service employment, wages, and rents explain only 11.8 percent of the variation in tracts' service-job gains. Home-route exposure alone explains 75.5 percent. Adding work-route exposure raises this to 76.7 percent, and adding the tract controls back moves it only to 77.6 percent. Where an entrant's workers live and shop largely determines which neighborhoods gain service jobs. Conditional on those factors, service jobs are lower in higher-wage places. Appendix E.5 defines the exposures and repeats the exercise leaving out each CZ in turn.

Table E.5: Commuting and shopping exposure and tract service-job changes

	(1)	(2)	(3)	(4)
Home-route exposure		0.84***	0.62***	0.59***
		(0.03)	(0.05)	(0.06)
Work-route exposure			0.45***	0.47***
			(0.11)	(0.11)
Log(1 + baseline service jobs)	0.05***			0.01***
	(0.01)			(0.00)
Log baseline service wage	0.01			−0.08***
	(0.01)			(0.01)
Log baseline residential rent	0.00			0.00
	(0.00)			(0.00)
Event × county effects	Yes	Yes	Yes	Yes
Within-county R^2 (%)	11.84	75.50	76.66	77.63
Tracts	114,782	114,782	114,782	114,782
CZ clusters	95	95	95	95

Notes: The outcome is the change in service jobs per 100 promised jobs. The expansion tract is excluded. Exposures are dollars reaching each tract per $100 of baseline payroll in the location-sector, including housing and service expenditure shares and the 0.74/0.26 home/work routing split. Commuting is weighted by skill-specific wages. Parentheses give CZ-clustered standard errors.

E.6 Standardized shocks and market characteristics

The 500-job and $70 million payroll designs each cover 216 observed sites in 95 CZs and both tradable industries, without a property-tax exemption. Productivity is chosen to add either 500 jobs or $70 million in annual tract–industry payroll, including wage changes on existing jobs. Both designs use the current labor calibration, $\kappa_N = 0.01$, CZ-specific $\sigma_{N,c}$, and no outside service-

shopping option. Figure 12 reports service jobs per 100 added jobs for the first design and per $10 million of additional payroll for the second. Per-job ratios use each run's realized employment increase. These are static comparisons, not transition paths.

Table E.6: Standardized shocks by entrant industry

Design / industry	Events	Service jobs	Net jobs	EV/resident	Payroll/job
500 jobs TH	216	29.49	80.40	6.49	105,576.03
500 jobs TL	216	17.99	71.35	4.33	75,511.85
$70 million TH	216	30.93	85.11	8.25	106,229.47
$70 million TL	216	20.49	76.09	8.42	77,783.40

Notes: Event-weighted medians. Service and net jobs are per 100 jobs added in the target location-sector. Welfare is annual dollars per baseline resident; payroll per job is annual incremental payroll in the target location-sector divided by its job increase. TH and TL denote high- and low-skill-intensive tradables. Sample sizes are identical across industries and designs: 216 events in each row. All dollars are in 2023 prices.

The cross-market regressions include industry effects and exclude the calibrated service expenditure share. Predictors are standardized after removing industry means. Shopping leakage is zero in every run, so its coefficient is not identified and is shown as a dash. Residential supply elasticity is weighted by baseline residents. The inbound share is inbound workers as a fraction of jobs in the CZ, not the share of residents working outside it. Standard errors use CZ-clustered CR1 covariance with the absorbed effects in the finite-sample adjustment; significance uses two-sided t_{94} tests. These are descriptive projections, not uncertainty intervals for the structural parameters or the numerical solutions.

Table E.7: Service multipliers and market characteristics: 500-job shock

	(1)	(2)	(3)	(4)	(5)	(6)
Log baseline employed residents	-1.55*	-2.26***	-0.38	-1.55*	-1.02	0.61
	(0.83)	(0.73)	(1.06)	(0.83)	(0.92)	(1.50)
Inbound employment share		-1.49**			-1.64**	-1.02
		(0.69)			(0.67)	(0.66)
Residential land supply elasticity			1.66*		1.86**	1.31*
			(0.85)		(0.80)	(0.73)
Shopping leakage				–	–	–
Shopping elasticity, $\sigma_{N,c}$						-0.57
						(1.57)
Commuting elasticity, χ_{ce}						-2.10*
						(1.07)
Industry effects	Yes	Yes	Yes	Yes	Yes	Yes
Within-industry R^2 (%)	3.60	6.18	5.67	3.60	8.76	13.44
Observations	432	432	432	432	432	432
CZ clusters	95	95	95	95	95	95

Notes: Same definitions as Table E.8. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table E.8 relates the service multiplier to characteristics of the market and of the site, all measured in the no-entry equilibrium. The outcome is the log of CZ service jobs per 100 added

jobs when employment in the entrant's industry at the site rises by 500. The sample is each observed site and its alternatives in the host county and in the other counties of the CZ, 4,392 sites in the 216 events. Six variables describe the CZ. Market size is the log of employed residents. The inbound employment share is the share of the CZ's jobs held by workers who live outside it. The residential land supply elasticity is the average across tracts, weighted by residents. The shopping elasticity $\sigma_{N,c}$ and the commuting elasticity χ_{ce} are the values used in each event's calibration, and they govern substitution across shopping destinations and across residence–workplace pairs. Service job concentration is the sum of squared tract shares of the CZ's service jobs. Three variables describe the site. Location-sector jobs are the jobs in the entrant's industry in the site's tract. The relative wage is the average wage of those jobs divided by the average wage across all jobs in the CZ. The inbound-commuting share is the share of those workers who live outside the CZ. Every column includes industry effects, and each coefficient is the change in the log multiplier for a one standard deviation change in the variable.

Table E.8: Market and site characteristics and the service multiplier

	(1)	(2)	(3)
Commuting-zone characteristics			
Log baseline employed residents	−0.02		0.00
	(0.06)		(0.04)
CZ inbound employment share	−0.05*		−0.01
	(0.03)		(0.01)
Residential land supply elasticity	0.03		0.04**
	(0.03)		(0.02)
Shopping elasticity, $\sigma_{N,c}$	0.02		0.01
	(0.06)		(0.03)
Commuting elasticity, χ_{ce}	−0.12***		−0.10***
	(0.04)		(0.02)
Service job concentration	−0.04		−0.03***
	(0.02)		(0.01)
Site characteristics			
Log location-sector jobs		−0.20***	−0.20***
		(0.01)	(0.01)
Log location-sector wage relative to CZ average wage		0.28***	0.27***
		(0.01)	(0.01)
Site inbound-commuting share		−0.07***	−0.06***
		(0.01)	(0.01)
Industry effects	Yes	Yes	Yes
R^2 after removing effects (%)	12.7	72.2	81.5
Sites	4,392	4,392	4,392
Events	216	216	216
CZ clusters	95	95	95

Notes: The outcome is the log of CZ service jobs per 100 added jobs under the 500-job shock. The sample is 4,392 sites in 216 events: each observed site and its alternatives in the host county and in the other counties of the CZ, with the industry held at its observed type. All columns include industry effects. Each coefficient is the change in the log multiplier associated with one standard deviation of the predictor after removing those effects. Service job concentration is the sum of squared tract shares of the CZ's baseline service jobs. The relative wage is the log of the location-sector wage minus the log of the CZ's average wage across all jobs at baseline. Parentheses give CZ-clustered standard errors; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, using two-sided t tests with one fewer degree of freedom than the number of CZ clusters.

F Counterfactual Costs, Locations, and Entrant Types

This appendix documents the current-model experiments in Section 7. They use the same 55-event labor calibration as the main results, $\kappa_N = 0.01$, CZ-specific $\sigma_{N,c}$, and no outside service-shopping destination for inside origins. The labor parameters are held fixed across experiments. Property-tax rates finance county service needs, with no government cash transfer to residents. Dollar values are annual 2023 dollars. Welfare is residential-opportunity equivalent variation and is evaluated separately for each geography.

F.1 Fiscal returns

The fiscal exercise contains 54 certified hypothetical entry scenarios, drawn from 70 records with verified local-benefit coverage. 15 records lack both a usable dollar amount and a documented percentage, and 1 lacks a supported baseline. None is assigned zero cost. The priced sample retains 3 duplicate references and treats retention-only, never-hired, and uncertain-site records as hypothetical expansions. It is a scenario summary, not an estimate for 70 independent realized awards. Existing baselines are reused; 33 priced records have revised hiring-year information and 1 has a revised site that is not installed in this exercise.

For 43 scenarios, let B_e be the annual local benefit in 2023 dollars. Sources include measured benefits, contractual amounts, and flagged monetary proxies. Source totals are divided by the documented term, with a ten-year default where the term is unavailable. Where annual observations are available, each is deflated using its own price year before averaging; missing years are not zeros. Two supplemental amounts use a property-level PILOT saving and a later-investment appropriation, respectively; neither is treated as verified event-specific realized savings. Remaining source and investment-vintage qualifications are retained in the underlying records.

Let K_e^{add} denote the new commercial capacity, measured using the baseline space requirement for the promised-job expansion. Let $j(e)$ and $a(e)$ denote its tract and county. Write $C_e(\xi)$ for the annual tax expenditure after solving the entry equilibrium with exemption fraction ξ, as defined in equation (F.4). I choose $\xi_e \in [0,1]$ to satisfy

$$C_e(\xi_e) = \min\{B_e, C_e(1)\}, \qquad \xi_e = 1 \quad \text{if } B_e > C_e(1). \tag{F.1}$$

The tax rate and rental base are recomputed in each candidate equilibrium. A full-exemption solve establishes the feasible upper amount; an iterative search then matches smaller recorded amounts within the greater of one dollar and one millionth of B_e. This matches 37 monetary-input scenarios and caps 6 at full exemption. For capped scenarios, the excess recorded amount remains in the input audit and is not paid in the model. For the 11 scenarios without dollar amounts, ξ_e is the documented percentage. These include conditional maxima, phased percentages averaged across their documented years, tax-linked grants, and personal-property exemptions applied as proxies to the model's commercial capacity. Their costs are model-implied, not observed tax savings.

Productivity is calibrated to the promised-job target without an exemption. That productivity and K_e^{add} remain fixed when the exemption is introduced; the labor parameters are not re-estimated. Commercial space has a common gross occupancy price, so the owner rent on exempted space is

$$r_e^{C,\text{new}} = \frac{1 + t_{a(e)}}{1 + (1 - \xi_e) t_{a(e)}} r_{j(e)}^{C}. \tag{F.2}$$

Writing $\mathcal{R}_{a(e)}^{C,\text{old}}$ for current rental income on nonexempt commercial space, the host county's budget becomes

$$t_{a(e)} \left[\mathcal{R}_{a(e)} + \mathcal{R}_{a(e)}^{C,\text{old}} + (1 - \xi_e) r_e^{C,\text{new}} K_e^{\text{add}} \right] = G_{a(e)}. \tag{F.3}$$

Service needs $G_{a(e)}$ follow equation (30). No government cash transfer is paid to residents. The annual tax expenditure is the revenue forgone on the exempt fraction of new capacity at equilibrium prices and the equilibrium county rate,

$$C_e(\xi_e) = \xi_e t_{a(e)} r_e^{C,\text{new}} K_e^{\text{add}}. \tag{F.4}$$

For an uncapped monetary-input scenario at its observed site, the solved cost equals B_e up to the matching tolerance. A capped scenario costs $C_e(1) < B_e$. In relocation experiments, the observed-site exemption fraction is held fixed, so the dollar cost changes with the site's equilibrium tax rate and rental base. Below, C_e abbreviates $C_e(\xi_e)$.

For geography A, cost per net job is $C_e/\Delta L_{eA}$, where ΔL_{eA} includes the expansion, changes at other tradable producers, and service jobs relative to no entry. Welfare per public dollar is EV_{eA}/C_e, using annual residential-opportunity equivalent variation. These quantities use the same cost numerator for the CZ and host county. Welfare is evaluated separately for each geography and is not additive across them. Each scenario receives equal weight; quantiles of ratios are not ratios of the separately reported quantiles. All 54 priced solves passed the equilibrium and export checks.

Table F.1: Local-benefit scenarios by input basis

Input group	Scenarios	Cost/job CZ	Cost/job County	Welfare/$ CZ	Welfare/$ County
Monetary inputs, uncapped	37	1,549.95	1,358.53	8.83	1.54
Monetary inputs, capped	6	7,738.57	5,941.76	1.26	0.50
Documented percentages	11	4,983.63	4,946.84	1.46	0.78

Notes: Within-group medians; annual 2023 dollars. Monetary inputs include flagged proxies. Capped cases receive a full exemption on added capacity; percentage-only cases use model-implied costs. The three groups partition the 54 priced scenarios. Welfare per public dollar is a conditional-on-entry ratio, not the causal return to the exemption.

On the same 43 monetary-input scenarios, dividing the full annual recorded amount B_e by net CZ jobs gives a median of $1,883.71, compared with $1,883.71 using the feasible modeled cost C_e. The pooled figure in Table 6 also includes the 11 percentage-only scenarios. No full state-and-local award or lifetime cost is inferred from the annual local-benefit amount.

Additional property-tax receipts finance service needs and are not credited as unrestricted funds against C_e. These ratios condition on entry and do not identify the subsidy-induced employment gain. The additionality calculation in Appendix F.6 compares offering the exemption with entry that would occur without it.

F.2 Alternative sites and paired payroll comparisons

Tables F.3 and F.4 use the current-model runs with no outside service shopping and no property-tax exemption. Each event retains its observed tradable industry. The candidate set contains the observed site and up to ten alternative tracts in its host county, inherited from the existing reproducible sample (seed 20261028). Unsupported sites remain unpriced and are not replaced.

Table F.2: Fiscal scenarios on the same supported baselines

Measure	No exemption	Recorded exemption
Commuting zone		
Service jobs per 100 promised jobs	29.72	27.67
Net jobs per 100 promised jobs	79.63	76.37
Annual welfare per baseline resident ($)	20.89	16.98
Annual tax expenditure per net job ($)	0.00	2,869.69
Resident welfare per public dollar	–	3.06
Host county		
Service jobs per 100 promised jobs	18.90	17.22
Net jobs per 100 promised jobs	93.44	90.73
Annual welfare per baseline resident ($)	31.88	24.99
Annual tax expenditure per net job ($)	0.00	2,575.37
Resident welfare per public dollar	–	1.09

Notes: Equal-weight medians for the same 54 scenarios under no exemption and the calibrated recorded exemption. Productivity shocks, added capacity, baselines, and labor parameters are held fixed. All employment and welfare changes are relative to no entry. Welfare per public dollar is undefined when tax expenditure is zero. Dollars are annual 2023 dollars. Differences between column medians are not medians of within-scenario differences.

The paired sample contains 2,250 supported sites in 214 events; two events without a supported county alternative are excluded. Both designs use exactly the same sites.

For the job design, productivity is chosen to add 500 tract–industry jobs at each site. For the payroll design, it is chosen to add $70 million in annual tract–industry payroll, including wage changes on existing jobs, with employment allowed to adjust. The industry and capacity rule are retained; the amount of added space may vary across locations. Income and welfare statistics use annual flows in 2023 dollars.

Let Y_{es}^{d} be an outcome at site s for event e under design $d \in \{J, P\}$. Table F.4 reports

$$\text{median}_e \left[\max_s Y_{es}^{d} - \min_s Y_{es}^{d} \right]$$

separately for each design and outcome. Job gaps are numbers of jobs, and welfare gaps are annual dollars per baseline resident. The best-site welfare percentage is the event median of $100[\max_s W_{es}^{d}/W_{eo}^{d} - 1]$, where W is total CZ welfare gain and o is the observed site under the same design. Sites maximizing different outcomes need not coincide. These comparisons describe evaluated locations, not an exhaustive county optimum.

Table F.3 uses the 500-job design. Its accounting decomposition uses within-event log deviations of the four income-transmission terms. The site regressions absorb event effects and standardize predictors after demeaning. Home- and work-route concentration are sums of squared shopping-destination shares. CZ-clustered CR1 standard errors and two-sided t_{93} tests use 94 CZ clusters.

F.3 County service gains and baseline site characteristics

Table F.5 uses the same 2,250 supported sites in 214 event grids and 94 CZs as Tables F.3 and F.4. Columns (1)–(3) use the 500-job design; the dependent variable is the host county's service-job

Table F.3: Decomposing the spread in the service multiplier across locations within a county

Panel A: Within-event decomposition of the log multiplier

Term	Share of within-event variance (%)	Median within-event P90/P10
Payroll per added job	55.53	1.36
Resident income per payroll dollar	54.53	1.38
CZ service spending per dollar of resident income	-13.17	1.10
Service jobs per dollar of CZ service spending	3.12	1.04
Service multiplier (product)	100.00	1.67

Panel B: Site characteristics and the within-event log multiplier

Site characteristic (baseline, standardized within event)	(1)	(2)	(3)
Log location-sector jobs	-0.18***		-0.17***
	(0.01)		(0.01)
Log location-sector wage	0.19***		0.19***
	(0.00)		(0.00)
Inbound-commuting share		-0.08***	-0.06***
		(0.01)	(0.00)
Home-route spending concentration		0.04***	0.01
		(0.01)	(0.01)
Work-route spending concentration		-0.05***	0.00
		(0.01)	(0.00)
Event effects	Yes	Yes	Yes
Within-event R^2 (%)	78.99	12.16	83.93
Locations	2,250	2,250	2,250
Events	214	214	214
CZ clusters	94	94	94

Notes: 214 events and 2,250 supported sites, including the observed sites and up to ten alternatives in each host county. The industry is held at its observed type, and each site receives 500 additional tract–industry jobs. Two events without a supported county alternative are excluded. The multiplier is CZ service jobs per 100 added jobs and equals 100 times the product of the four terms in Panel A. Payroll is the additional tract–industry wage bill; resident income is additional CZ disposable income; service spending is additional CZ service revenue. Contributions are covariances of within-event log deviations with the log multiplier, divided by its within-event variance. Panel B includes event effects; predictors are standardized after demeaning. Shopping concentration is the sum of squared destination shares for each route. Parentheses give CZ-clustered standard errors; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, using two-sided t tests with one fewer degree of freedom than the number of CZ clusters.

change per 100 added jobs. Columns (4)–(5) use the $70 million payroll design and report county service jobs per $10 million of additional payroll. Each regression absorbs event effects and gives every event equal total weight.

Home- and work-route exposures measure baseline dollars reaching county service businesses per $100 of gross payroll in the target location-sector. They combine skill-specific wages, commuting, housing and service expenditure shares, the 0.74/0.26 home/work origin split, and the inverted shopping probabilities. Controls are the log baseline mean wage, log one plus target location-sector employment, log one plus service employment in the tract, and the tract's residential supply elasticity. The calibrated service expenditure share enters the exposure accounting but is not an

Table F.4: Relocating an expansion at equal jobs or equal payroll

	500 additional jobs	$70 million payroll
Gap between the best and worst site		
Service jobs	111.00	63.96
Net jobs	81.89	187.86
Welfare per resident ($)	3.77	1.74
Best site compared with the observed site		
Welfare gain (%)	37.92	8.49

Notes: Both designs use the same 2,250 supported sites in 214 events, including each observed site and up to ten alternatives in its host county, with the industry held at its observed type. Entries are medians of within-event comparisons. All outcomes refer to the CZ. The first design fixes the target tract–industry employment increase at 500 jobs. The second fixes its additional annual payroll at $70 million (2023 dollars), including wage changes on existing jobs, and allows the employment increase to range from 198 to 1,399 jobs. Each outcome's best and worst sites are selected separately; the welfare percentage compares the welfare-maximizing site with the observed site under the same design. All runs have no property-tax exemption.

additional regressor. Required payroll in column (3) is annual incremental payroll per 100 added jobs, in millions of dollars. It is an equilibrium outcome, so this column is an accounting diagnostic rather than a causal estimate. Predictors retain their original units and are not standardized.

Table F.5: County service-job gains across alternative sites

	500 jobs (1)	500 jobs (2)	500 jobs (3)	$70m payroll (4)	$70m payroll (5)
Home-route exposure	0.73***	0.47***	0.55***	0.52***	0.48***
	(0.08)	(0.06)	(0.05)	(0.03)	(0.03)
Work-route exposure	0.39***	0.95***	0.71***	0.70***	0.78***
	(0.12)	(0.10)	(0.10)	(0.04)	(0.05)
Log baseline target location-sector wage		23.14***			4.52***
		(2.54)			(0.82)
Log(1 + target location-sector jobs)		-0.80***			-0.28***
		(0.30)			(0.07)
Log(1 + site service jobs)		-1.47***			-0.08***
		(0.12)			(0.03)
Residential supply elasticity		-0.47			0.41
		(1.31)			(0.36)
Required payroll (million dollars)			2.04***		
			(0.47)		
Within-event R^2 (%)	15.27	57.35	60.48	74.56	82.64
Observations	2,250	2,250	2,250	2,250	2,250
Events	214	214	214	214	214
CZ clusters	94	94	94	94	94
Event effects	Yes	Yes	Yes	Yes	Yes

Notes: The same complete paired sites enter every column. Columns (1)–(3) hold the added jobs at 500 and report county service jobs per 100 added jobs. Columns (4)–(5) hold incremental payroll at $70 million and report county service jobs per $10 million. Event fixed effects and equal total event weights are used throughout. Parentheses contain CZ-clustered CR1 standard errors; $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$, from two-sided t_{93} tests. Reported R^2 is conditional on event effects. These are descriptive relationships among simulated outcomes.

F.4 Recorded-benefit location comparisons

The fiscal relocation exercise covers 537 supported locations across the 54 priced scenarios, including each observed site and its supported host-county alternatives. Each scenario retains its original industry and promised-job target. At each site, productivity is calibrated without an exemption, then held fixed when applying the exemption fraction recovered at the observed site. Added commercial capacity follows the same rule and is also held fixed between that site's no-exemption and exemption solves. The exemption fraction is constant across sites, but the annual dollar cost changes with equilibrium rents and the property-tax rate. Cost per net job and cost per service job divide this site-specific cost by the respective entry-minus-no-entry employment change. A site with nonpositive service-job gains has no positive cost-per-service-job ratio and is excluded from that ratio comparison. The sampled locations do not establish an optimum over all tracts in the county.

F.5 Industry comparisons

Each of the 216 observed sites has expansions in both high- and low-skill-intensive tradables under the common 500-job and $70 million incremental-payroll designs. Both industries use the same site and baseline, with no property-tax exemption. Table E.6 reports the outcomes by industry; Figure 12 shows their distribution across local economies. The median within-event ratio of service jobs under high- versus low-skill-intensive entry is 1.55 with 500 added jobs and 1.05 with $70 million of additional payroll. The ratios are computed within events before taking medians. These comparisons hold either employment or payroll fixed, rather than holding the productivity shock or public expenditure fixed.

F.6 The additionality calculation

Each of the 54 recorded-subsidy scenarios in Table 6 is solved three times: with the recorded exemption, without it, and without entry. Let W_{ex} and W_{no} be residents' annual welfare relative to no entry when the firm enters with and without the exemption. If the exemption causes a share s of the entries it subsidizes, offering it yields W_{ex} in every scenario, while withholding it yields W_{no} from the share $1 - s$ of firms that would enter anyway and nothing from the rest. The expected gain from offering it is $W_{\text{ex}} - (1 - s) \cdot W_{\text{no}}$, which Table 7 divides by the exemption's annual tax expenditure. The exemption pays off when s exceeds $1 - W_{\text{ex}}/W_{\text{no}}$. In the median scenario, this threshold is 11.64 percent for CZ residents and 13.55 percent for host-county residents. Expected net jobs are computed the same way from net workplace jobs.