

Passenger Travel Costs and the Spatial Organization of Economic Activity

Zeyu Chen Deyu Rao Lin Yang Wen Wang Yatang Lin *

October 12, 2026

Abstract

We develop and estimate a quantitative spatial model to study how passenger connectivity affects production, migration, and welfare. Firms optimally adjust destination-specific supplier search intensity as travel costs change, while households value access to nonlocal leisure amenities. Combining card-transaction-based measures of intercity business and leisure travel with city-pair data on trade and migration, we quantify the effects of China's high-speed rail expansion between 2008 and 2015. Counterfactuals imply that the 2015 network raises national average real wages by 2.19% and welfare by 3.14% relative to a no-HSR baseline. Both channels contribute substantially to real-wage gains, whereas improved leisure accessibility accounts for most welfare gains. Uneven accessibility improvements concentrate economic activity in large connected cities and generate adverse spillovers on cities outside the network. Simulation for alternative network expansions reveals that real-wage and welfare gains are concentrated in high-volume links, while further expansion could reduce induced migration and spatial inequality in economic outputs.

*Chen: Department of Geography and Environment, London School of Economics and Political Science; Rao, and Lin: Department of Economics, Hong Kong University of Science and Technology; Yang: Society Hub, Hong Kong University of Science and Technology (Guangzhou); Wang: Division of Social Science, Hong Kong University of Science and Technology. Emails: z.chen144@lse.ac.uk, dyrao@ust.hk, yangl@ust.hk, wenwangww@ust.hk, linyt@ust.hk. We thank Shanjun Li, Binglin Wang, Wei Lin, Alberto Salvo, Zhi Wang, Jordan Norris, Jonathan Hall, Matthias Hoelzlein, Yuta Suzuki, Yuyang Jiang, Chen Zhang, Vernon Henderson, Giorgio Pietrabissa, and participants at Fudan CCES Urban & Spatial Workshop, HKUST(GZ) Environmental Workshop, Jinan GLO, HKUST-Fudan-SMU Trade Workshop, SEU Urban & Labor Workshop, for helpful comments, and Zhencheng Zhang and Jingchao Hu for excellent research assistance.

1 Introduction

Transportation infrastructure shapes the spatial organization of economic activity by determining how easily goods and people move across places. A large literature studies this question through freight costs and goods-market access. Yet investments such as high-speed rail (HSR) primarily reduce the cost of moving people rather than goods. Lower passenger travel costs affect two distinct sets of economic decisions. Firms travel to discover and evaluate suppliers, negotiate terms, and sustain trading relationships (Cristea, 2011; Startz, 2021). Households travel to consume amenities outside their home city, and lower passenger travel costs improve access to those amenities from a given residence (Faber and Gaubert, 2019; Albert and Monras, 2022). These movements are temporary, but they can have lasting effects on production networks, migration, local prices, and welfare.

Despite its importance, passenger mobility has not been fully incorporated into quantitative models of the spatial economy. On the production side, standard trade models treat bilateral trade costs and productivity distributions as primitives, leaving little role for the face-to-face interactions through which buyers discover suppliers and improve procurement relationships. On the consumption side, conventional measures of local amenities omit the value of reaching nonlocal destinations without changing residence. Existing work has separately emphasized business travel, tourism, migration, commuting, or the spatial reach of firms. We therefore lack a unified quantitative account of how a passenger network propagates through firm search, leisure travel, trade, migration, and welfare.

This paper addresses this gap by developing and estimating a quantitative spatial model in which passenger travel costs enter both firms' sourcing decisions and household utility. We use the model to study China's HSR expansion between 2008 and 2015. The framework introduces two channels. Through a *business-search channel*, final-goods producers choose where and how intensively to search for suppliers, with search costs increasing in passenger travel time. Through a *leisure-accessibility channel*, residents value access to destinations outside their residence city through short-term travel. Together, the channels provide a tractable account of how moving people changes both the production opportunities available to firms and the consumption opportunities available from a location.

The business-search channel embeds endogenous supplier search in a multi-sector Eaton–Kortum environment. A buyer chooses a continuous search intensity for every potential supplier city. More intensive search increases the number of suppliers encountered and improves the distribution of input prices available to the buyer, but requires labor and travel. A reduction in passenger travel time therefore induces the buyer to search more intensively in the affected destination, expand its effective supplier network, and obtain better input prices. Because search is destination-specific, the model generates bilateral business travel and trade shares and permits these two objects to discipline the search technology. Endogenous search also amplifies standard trade-cost shocks: lower sourcing costs reduce final-goods prices, which feed back into production costs, firm entry, real wages, and the spatial allocation of labor.

The leisure-accessibility channel distinguishes short-term travel from long-term migration.

Workers choose where to live and work based on real wages, local amenities, migration costs, and the expected value of leisure trips available from each residence. Conditional on where they live, they choose among nonlocal destinations that differ in attractiveness and travel costs. Aggregating these destination choices yields a travel-time-weighted index of nonlocal leisure accessibility. Two cities with similar wages and local amenities can thus offer different utility when one provides better access to surrounding destinations. Passenger connectivity affects welfare directly, by expanding residents' effective consumption opportunities, and indirectly, by changing relative residential attractiveness and inducing migration.

China's HSR expansion provides a useful setting in which to quantify these mechanisms. The network produced large and heterogeneous reductions in passenger travel time across more than 300 cities. The number of directed city pairs connected by HSR increased from 293 in 2008 to 6,669 in 2015, and the connected pairs' share of baseline national business travel rose from 16.2% to 66.7%. HSR carried passengers rather than freight, helping us distinguish travel-mediated supplier search from changes in the physical cost of shipping goods. Moreover, the early network was concentrated along high-volume corridors and among large eastern cities. This uneven expansion provides variation in bilateral treatment intensity and makes the topology of the network central to the distribution of gains.

Our analysis combines several datasets at the city-pair level. We construct annual door-to-door passenger travel times by integrating official HSR and conventional-rail schedules with a digitized highway network and local access times between city centers and railway stations. We measure bilateral business and leisure trips using nationwide UnionPay credit- and debit-card transactions. A cardholder's most frequently used city in a rolling twelve-month window identifies residence; weekday trips outside public holidays measure business travel, while trips during public holidays and the extended Chinese New Year period measure leisure travel. We combine these mobility data with bilateral migration flows from the 2010 Population Census and the 2015 one-percent Population Sample Survey, and with city-level input-output tables that measure trade across 42 sectors in 2012 and 2015. The counterfactual analysis covers 302 cities for which all model inputs are available.

We first estimate leisure destination choice using a gravity equation. A 1% reduction in passenger travel time raises a destination's share of leisure trips by about 2%. The leisure accessibility recovered from origin fixed effects then enters the residential-choice equation, where we instrument it using geographic proximity to attractive external destinations. The estimates imply that the elasticity of residential choice with respect to leisure accessibility is about 9.5% of the calibrated elasticity with respect to real wages.

On the production side, business travel and trade jointly discipline the search technology. Conditional on the sourcing demand implied by the model, a 1% reduction in passenger travel time raises bilateral business travel by about 0.3%. We then estimate the trade gravity equation, and our preferred specification implies that a 1% reduction in passenger travel time raises bilateral trade shares by about 0.22%. Combined with the calibrated freight-cost elasticity, these estimates imply that the trade elasticity with respect to passenger travel costs is about 12% of that with respect to freight transport costs.

We use the estimated model to quantify the long-run effects of the historical HSR travel-time shocks. Starting from a recovered no-HSR equilibrium, we introduce the travel-time reductions associated with the network observed in each year while holding underlying fundamentals fixed. By 2015, the model implies that HSR lowers national average business travel time by 18.4%; its direct effect lowers national average leisure travel time by 15.7%. These reductions are highly uneven: among connected cities, the direct effects are substantially larger in large cities than in small cities for both types of travel. This heterogeneity also underscores why a binary indicator for HSR access is insufficient.

These travel-time reductions generate economically meaningful aggregate gains. Relative to the no-HSR baseline, the 2015 network raises national average real wages by approximately 2.19% and national average welfare by approximately 3.14%. Supplier search raises productivity by improving access to inputs, while leisure accessibility raises utility directly and changes the allocation of workers across cities. The direct leisure-accessibility gain accounts for the majority of the welfare effect, while both channels contribute to the real-wage gain. Using 2015 GDP as a rough monetary benchmark, one year of welfare-equivalent gains is strikingly close to the cumulative construction cost of the HSR network in place by 2015: the ratio of the former to the latter ranges from approximately 0.88 to 1.31. Passenger infrastructure can therefore affect production without carrying freight and generate welfare gains beyond those captured by conventional measures of goods-market access.

The same counterfactual also produces substantial population reallocation. This response does not arise because HSR directly lowers migration costs, which remain fixed in the model. Instead, uneven changes in real wages and nonlocal leisure accessibility alter cities' relative attractiveness. Relative to the no-HSR baseline, the simulated 2015 network increases the population of large connected cities by approximately 6.81 million and reduces the population of cities without HSR by approximately 6.94 million, while leaving small connected cities nearly unchanged. Most of the gain in large connected cities comes from workers whose *hukou* origins are in cities outside the network rather than in small connected cities. At the same time, the near-zero population change in small connected cities conceals offsetting inflows and outflows. The model-implied distributional pattern is therefore not a simple core-periphery siphoning story; it reflects the intensity and topology of cities' exposure to the passenger network.

Counterfactual expansions beyond the 2015 network reveal a further distinction between welfare gains and spatial reallocation. Because migration responds to relative rather than absolute accessibility, the migration response need not increase monotonically with network coverage. Broader coverage continues to raise welfare through improved access to nonlocal amenities but can reduce migration as accessibility gains become more evenly distributed across cities. Welfare gains from passenger connectivity therefore need not be accompanied by greater spatial reallocation.

Our study contributes to three strands of the literature. First, we embed firms' endogenous choice of destination-specific search intensity in a tractable quantitative spatial model. A growing body of research shows that information and matching frictions shape trade and the architecture of buyer-supplier relationships (Allen, 2014; Oberfield, 2018; Miyauchi, 2024), while some studies link face-to-face contact and passenger connections to trade, investment, and firm organization

(Cristea, 2011; Giroud, 2013; Gokan et al., 2019; Koster et al., 2021; Startz, 2021; Gumpert et al., 2022; Antràs et al., 2023). More closely related, Bernard et al. (2019) link passenger connectivity to firm performance through endogenous supplier search, and Eaton et al. (2026) and Hu and Xu (2024) develop quantitative frameworks that incorporate frictional input sourcing and buyer–seller matching into general equilibrium. Building on this work, firms in our model optimally adjust their search intensity in each location in response to passenger travel costs and the returns to sourcing. This margin provides a microfoundation for the passenger-travel component of trade costs and allows us to quantify how search responses generate spatial general equilibrium effects.

Second, we contribute to quantitative spatial economics literature (e.g., Allen and Arkolakis, 2014; Ahlfeldt et al., 2015; Caliendo and Parro, 2015; Tombe and Zhu, 2019) by separately estimating how intercity business and leisure travel respond to passenger travel time and integrating both channels into a unified spatial model. More closely related quantitative spatial models incorporate travel to consume nontradable goods and services and link consumption access to location choices (Rao, 2022; Huang et al., 2025; Miyauchi et al., 2025; Lee and Tan, 2026). The leisure channel also connects our framework to research on tourism and spatial development (Faber and Gaubert, 2019) and consumption outside the place of residence and location choice (Albert and Monras, 2022). Building on this literature, we use calendar timing in transaction data to construct measures for bilateral business and leisure travel volumes and estimate channel-specific responses to passenger travel time. We integrate firms’ supplier search and households’ access to nonlocal leisure amenities in a common spatial model, allowing us to quantify how the two channels’ contributions to real wages and welfare change as the passenger network expands.

Third, we contribute to research on transportation infrastructure and market access (e.g., Faber, 2014; Redding and Turner, 2015; Donaldson and Hornbeck, 2016; Donaldson, 2018; Ma and Tang, 2024) and, more specifically, on HSR and passenger connections (Lin, 2017; Qin, 2017; Ahlfeldt and Feddersen, 2018; Campante and Yanagizawa-Drott, 2018; Koster et al., 2022; Koster and Thisse, 2026), by quantifying how heterogeneous improvements in passenger accessibility and spatial spillovers shape aggregate benefits and local outcomes. Complementing empirical evidence on the local effects of HSR openings, we characterize cities’ direct exposure through bilateral travel-time reductions and the importance of the affected destinations to local firms and residents. We then trace how these changes propagate through trade and migration, accounting for adjustments in connected cities and spillovers on cities outside the network.

The remainder of the paper is organized as follows. Section 2 develops the quantitative spatial model with endogenous business search and leisure travel. Section 3 describes the construction of our datasets and documents the evolution of China’s HSR network. Section 4 presents the calibration and structural estimation. Section 5 quantifies the aggregate effects of HSR network construction on travel times, the spatial allocation of labor, real wages, and welfare. Section 6 examines the sources of heterogeneous effects across cities. Section 7 compares the historical network with a universal network and with rollouts that add links in descending order of baseline travel volume. Section 8 discusses the broader implications for infrastructure policy, and Section 9 concludes.

2 A quantitative spatial model with business and leisure travel

We develop a quantitative spatial model to analyze how the HSR network reshapes the spatial organization of economic activity by simultaneously reducing business and leisure travel frictions. Our framework builds upon the seminal contributions in spatial economics (Arkolakis, 2010; Ahlfeldt et al., 2015; Caliendo and Parro, 2015; Miyauchi et al., 2025) but departs from them by explicitly embedding HSR-driven accessibility into two channels of spatial equilibrium: First, we introduce an intercity leisure travel decision into individuals' sequential location choices. We conceptualize nonlocal leisure accessibility as a distinct, network-driven locational attribute that yields direct utility. As HSR is deployed unevenly across space, it generates asymmetric improvements in this leisure accessibility, thereby shifting the spatial distribution of labor. Second, we allow business travel frictions to shape the effective cost of sourcing intermediate inputs across space. When passenger travel between two cities is costly, final-goods buyers find it more expensive to evaluate suppliers, verify quality, negotiate contracts, and sustain procurement relationships in the distant market; these informational frictions raise the effective import prices that buyers can achieve. The HSR network, by lowering bilateral business travel barriers, therefore compresses effective sourcing costs on connected routes and restructures spatial trade patterns. We microfound this mechanism with an endogenous search technology: more intensive search improves the distribution of import price offers that buyers encounter from each origin, and firms face the trade-off between this cost reduction and the travel-time-dependent search costs. Search thus provides a tractable way to discipline how business travel time enters sourcing costs and equilibrium trade shares.

The spatial economy comprises N discrete locations, indexed by $j, n, m \in \mathcal{N} \equiv \{1, \dots, N\}$, and K production sectors, indexed by $k, s \in \mathcal{K} \equiv \{1, \dots, K\}$. The general equilibrium operates through three interacting blocks: (i) *Workers* endogenously choose their residential location, supply labor, consume goods and housing, and make intercity leisure travel decisions; (ii) *Intermediate goods firms* operate under perfect competition, utilizing labor and final goods to produce sector-specific varieties; (iii) *Final goods firms* operate under monopolistic competition with a fixed entry cost. They aggregate intermediate varieties into non-tradable final goods for local consumption and production. Crucially, their sourcing choices and thereby the import prices they can achieve are constrained by costly search over intermediate suppliers.

2.1 Workers

We assume that an individual i , who originates from city j , resides and works in city n , and travels to city m ($m \neq n$) for leisure purposes, obtains the following utility:

$$u_{jnm}(i) = \frac{a_n \varepsilon_n(i)}{\iota_{jn}} \left(\frac{C_n(i)}{\beta} \right)^\beta \left(\frac{H_n(i)}{1-\beta} \right)^{1-\beta} \left(\frac{b_m \xi_m(i)}{\mu_{nm}} \right)^\psi, \quad j, n \in \mathcal{N}, \quad m \in \mathcal{N} \setminus \{n\}.$$

In this specification, a_n denotes the average residential amenity of city n , encompassing local public goods and intra-city leisure facilities. By explicitly imposing the restriction $m \neq n$, we analytically separate local amenities from nonlocal leisure travel, isolating the unique welfare gains derived

from accessing external locations. The term $\varepsilon_n(i)$ represents an idiosyncratic preference shock for residing in city n , drawn independently from a Fréchet distribution with cumulative distribution function $G_\varepsilon(x) = \exp(-x^{-\theta^R})$. The parameter ι_{jn} captures the bilateral migration frictions from origin j to destination n .

Workers derive utility from the consumption of final goods, $C_n(i)$, and housing, $H_n(i)$. The budget constraint is given by:

$$p_n C_n(i) + r_n H_n(i) = w_n,$$

where p_n is the aggregate price index for final goods, r_n is the housing rent, and w_n is the nominal wage rate in city n . The Cobb–Douglas preference implies constant expenditure shares, with $p_n C_n(i) = \beta w_n$ and $r_n H_n(i) = (1 - \beta) w_n$.

Final goods consumption is structurally nested. In city n , an endogenous number of F_n^k final goods firms operate in sector k , indexed by $f \in \mathcal{F}_n^k \equiv \{1, \dots, F_n^k\}$. The aggregate final goods consumption is given by:

$$C_n(i) = \prod_{k=1}^K \left[\frac{1}{\rho_n^k} \left(\sum_{f=1}^{F_n^k} C_n^{k,f}(i)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \right]^{\rho_n^k}, \quad \sigma > 1, \quad \sum_{k=1}^K \rho_n^k = 1,$$

where $C_n^{k,f}(i)$ is the consumption of the final good produced by firm f in sector k of city n . Goods within a sector are substitutable with a constant elasticity of substitution $\sigma > 1$, while sector-specific composites are aggregated via a Cobb–Douglas function with expenditure shares ρ_n^k .

Utility derived from leisure travel is determined by three components: the intrinsic destination attractiveness b_m , the leisure travel friction μ_{nm} , and an idiosyncratic destination-specific preference shock $\xi_m(i)$ drawn from a Fréchet distribution $G_\xi(x) = \exp(-x^{-\theta^T})$. We model the travel friction μ_{nm} purely as an iceberg-style utility discount rather than a monetary cost in the budget constraint. This microfoundation captures the time-intensive nature of travel: time spent in transit directly crowds out the effective duration one can enjoy the destination's leisure facilities, thereby proportionally discounting the realized utility. Furthermore, unlike models of high-frequency intra-metropolitan trips where individuals allocate monetary budgets across various destinations (Miyauchi et al., 2025), our context of HSR-driven intercity leisure travel represents a relatively low-frequency activity. Abstracting from the monetary expenditures of leisure travel and a dedicated tourism sector allows us to maintain analytical tractability.

Following the sequential optimization framework in Miyauchi et al. (2025) and Lee and Tan (2026), we assume that workers make their location, travel, and consumption decisions sequentially, with idiosyncratic preference shocks revealed in stages. The timeline of the worker's problem unfolds as follows:¹

Stage 1. *Residential choice.* Upon observing the realization of the amenity preference shocks $\{\varepsilon_n(i)\}_{n \in \mathcal{N}}$, the worker selects an optimal residential and working city n . Crucially,

¹Because the leisure travel and consumption decisions are independent of one another, Stages 2 and 3 can alternatively be interpreted as simultaneous decisions.

at this stage, the destination-specific travel preference shocks $\{\xi_m(i)\}_{m \in \mathcal{N} \setminus \{n\}}$ remain unobserved; hence, the residential choice is based on the *expected* utility of future travel.

Stage 2. *Leisure travel choice.* Conditional on the chosen residential city n , the worker observes the travel preference shocks $\{\xi_m(i)\}_{m \in \mathcal{N} \setminus \{n\}}$ and selects an optimal leisure travel destination m to maximize travel utility.

Stage 3. *Consumption and housing allocation.* Given the chosen residential city, the worker allocates their wage income between housing and final goods.

We solve this sequential optimization problem via backward induction, starting from the final stage. Given prices and wages in city n , individual demand for a specific final good from firm f can be shown to be

$$C_n^{k,f}(i) = \frac{(p_n^{k,f})^{-\sigma}}{(p_n^k)^{1-\sigma}} \rho_n^k \beta w_n, \quad (1)$$

where $p_n^{k,f}$ is the firm-specific price, and

$$p_n^k \equiv \left[\sum_{f=1}^{F_n^k} (p_n^{k,f})^{1-\sigma} \right]^{\frac{1}{1-\sigma}} \quad (2)$$

is the sector-specific price index. Substituting Equation (1) back into the utility function yields the indirect utility conditional on residential location choices:

$$u_{jnm}(i) = \frac{a_n \varepsilon_n(i)}{\iota_{jn}} \frac{w_n}{P_n} \left(\frac{b_m \xi_m(i)}{\mu_{nm}} \right)^\psi,$$

where the aggregate price index for final goods is $p_n \equiv \prod_{k=1}^K (p_n^k)^{\rho_n^k}$, and $P_n \equiv p_n^\beta r_n^{1-\beta}$ represents the local consumer price index.

Moving backward to Stage 2, conditional on residing in city n and upon observing the realization of the travel preference shocks $\{\xi_m(i)\}_{m \in \mathcal{N} \setminus \{n\}}$, workers choose their optimal leisure travel destination to maximize $b_m \xi_m(i) / \mu_{nm}$. Integrating over the Fréchet distribution of these idiosyncratic shocks, the ex-post leisure travel share of residents in city n who travel to city m perfectly coincides with the ex-ante choice probability:

$$\lambda_{m|n}^T = \Pr \left\{ \frac{b_m \xi_m(i)}{\mu_{nm}} \geq \max_{m' \in \mathcal{N} \setminus \{n\}} \left\{ \frac{b_{m'} \xi_{m'}(i)}{\mu_{nm'}} \right\} \right\} = \frac{(b_m / \mu_{nm})^{\theta^T}}{\Phi_n^T}, \quad m \in \mathcal{N} \setminus \{n\}, \quad (3)$$

where $\Phi_n^T \equiv \sum_{m' \in \mathcal{N} \setminus \{n\}} (b_{m'} / \mu_{nm'})^{\theta^T}$.

In Stage 1, individuals do not possess information on the exact realization of $\{\xi_m(i)\}$ when selecting their residential location. Instead, they evaluate the expected utility of residing in city n

based on the distribution of travel shocks:

$$U_{jn}(i) \equiv \mathbb{E}_{\xi} \left[\max_{m \in \mathcal{N} \setminus \{n\}} \left\{ \frac{a_n \varepsilon_n(i)}{\iota_{jn}} \frac{w_n}{P_n} \left(\frac{b_m \xi_m(i)}{\mu_{nm}} \right)^{\psi} \right\} \right] = \frac{a_n \varepsilon_n(i)}{\iota_{jn}} \frac{w_n}{P_n} \bar{\gamma} \left(\Phi_n^T \right)^{\psi/\theta^T}, \quad (4)$$

where $\bar{\gamma} \equiv \Gamma(1-\psi/\theta^T)$ is a constant.² Constructed precisely as a travel-cost-weighted sum of external destination attractiveness, Φ_n^T forms a quantitative index for the overall leisure travel accessibility of city n . As shown in Equation (4), this network-driven accessibility acts as a fundamental locational attribute. It enters the residential decision symmetrically to local real wages and intra-city amenities, formally capturing the spatial premium of being well-connected to external leisure destinations.

Importantly, the composite elasticity ψ/θ^T governs the relative importance of this “outside” leisure accessibility versus the “inside” local utility. Encapsulating accessibility via this expected-utility framework allows us to remain agnostic about individual travel frequencies and the extensive margin of non-travelers. We do not need to explicitly model a subset of residents who never travel or pinpoint the exact number of trips per year. If intercity leisure travel is inherently low-frequency or weakly valued, the preference parameter ψ will be small. Furthermore, if a substantial proportion of residents rarely travel out, which implies that they are highly insensitive to destination-specific idiosyncratic shocks, the Fréchet shape parameter θ^T will be correspondingly large. In either scenario, the structural term ψ/θ^T naturally compresses toward zero. We defer the determination of the leisure channel’s aggregate importance to the data via the subsequent parameter estimation.

Upon observing the residential preference shocks $\{\varepsilon_n(i)\}_{n \in \mathcal{N}}$, workers optimally choose their residence $n^*(i) = \arg \max_{n \in \mathcal{N}} U_{jn}(i)$. The equilibrium share of workers migrating from origin j to destination n is thus given by:

$$\lambda_{jn} = \Pr \left\{ U_{jn}(i) \geq \max_{n' \in \mathcal{N}} \{U_{jn'}(i)\} \right\} = \frac{(a_n w_n)^{\theta^R} (\iota_{jn} P_n)^{-\theta^R} (\Phi_n^T)^{\theta^R \psi/\theta^T}}{\Theta_j}, \quad (5)$$

where $\Theta_j \equiv \sum_{n'=1}^N (a_{n'} w_{n'})^{\theta^R} (\iota_{jn'} P_{n'})^{-\theta^R} (\Phi_{n'}^T)^{\theta^R \psi/\theta^T}$. Therefore, the total labor supply in city n is given by

$$L_n = \sum_{j=1}^N \lambda_{jn} \bar{L}_j,$$

where $\bar{L}_j$ is the exogenous initial labor endowment in origin city j .

Finally, we aggregate the individual demand for final goods. While workers allocate a share $1-\beta$ of their wage income to housing, we assume that all housing rents are collected by local landlords, who share the exact same preference structure as workers and spend their entire housing income on local final goods. Consequently, the aggregate expenditure on final goods in city n precisely equals the total local wage bill, $w_n L_n$. Aggregating the individual demand from Equation (1) across all workers and landlords, the aggregated consumption demand faced by final goods firm f in sector k

²We use $\Gamma(\cdot)$ to denote the Gamma function, thus this requires $\psi/\theta^T < 1$. Our subsequent parameter estimation will verify this quantitative assumption.

of city n is given by

$$C_n^{k,f} = \frac{(p_n^{k,f})^{-\sigma}}{(p_n^k)^{1-\sigma}} \rho_n^k w_n L_n. \quad (6)$$

2.2 Firms

2.2.1 Intermediate goods production

In each sector k and city j , intermediate goods are produced by perfectly competitive firms. We conceptualize the product space as a continuum of sector-specific intermediate varieties indexed by $\omega^k \in [0, 1]$. For any specific variety ω^k , there exists a pool of potential intermediate suppliers in city j , each endowed with an idiosyncratic productivity draw $z_j^k(\omega^k)$.

To produce variety ω^k , an intermediate firm utilizes two types of inputs: labor and final goods (also referred to as materials) from all sectors. The production function of a firm with productivity $z_j^k(\omega^k)$ is given by

$$y_j^k(\omega^k) = z_j^k(\omega^k) \left[L_j^k(\omega^k) \right]^{\gamma_j^k} \prod_{s=1}^K \left[q_j^{k,s}(\omega^k) \right]^{\gamma_j^{k,s}}, \quad \gamma_j^k + \sum_{s=1}^K \gamma_j^{k,s} = 1,$$

where $L_j^k(\omega^k)$ is the labor input, and $q_j^{k,s}(\omega^k)$ is the input of sector s 's final goods produced in city j , defined by

$$q_j^{k,s}(\omega^k) = \left[\sum_{f=1}^{F_j^s} q_j^{k,s,f}(\omega^k)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}},$$

where $q_j^{k,s,f}(\omega^k)$ is the input of sector s 's final good produced by firm f for variety ω^k . Note that this formulation shares the exact same CES aggregator as final goods consumption; thus, the aggregate input cost of sector-specific materials is equal to the sector-specific final goods price index p_j^s defined previously.

Due to the homotheticity of the Cobb–Douglas production function, the unit cost of the input bundle (prior to the productivity adjustment) is given by

$$c_j^k = \bar{\gamma}_j^k w_j^{\gamma_j^k} \prod_{s=1}^K (p_j^s)^{\gamma_j^{k,s}},$$

where $\bar{\gamma}_j^k \equiv (\gamma_j^k)^{-\gamma_j^k} \prod_{s=1}^K (\gamma_j^{k,s})^{-\gamma_j^{k,s}}$ is a city-sector-specific constant. Consequently, the marginal cost for a specific intermediate firm to produce variety ω^k is $c_j^k / z_j^k(\omega^k)$. Furthermore, cost minimization implies that the intermediate firm's conditional demand for the final good produced by firm f in sector s can be derived as:

$$q_j^{k,s,f}(\omega^k) = \frac{(p_j^{s,f})^{-\sigma}}{(p_j^s)^{1-\sigma}} \frac{\gamma_j^{k,s}}{\gamma_j^k} w_j L_j^k(\omega^k).$$

Integrating over the continuum of varieties $\omega^k \in [0, 1]$ and summing across all demanding sectors $k \in \mathcal{K}$, the aggregate intermediate demand faced by final goods firm f in sector s of city j is given by

$$Q_j^{s,f} = \sum_{k=1}^K \frac{(p_j^{s,f})^{-\sigma}}{(p_j^s)^{1-\sigma}} \frac{\gamma_j^{k,s}}{\gamma_j^k} w_j L_j^k, \quad (7)$$

where L_j^k denotes the total production labor employed by intermediate firms in sector k of city j .

2.2.2 Final goods production and search frictions

In each city n and sector k , an endogenous mass of final goods firms, indexed by $f \in \mathcal{F}_n^k$, operates under monopolistic competition. Each firm f acts as an aggregator, sourcing the continuum of intermediate varieties $\omega^k \in [0, 1]$ across various locations to produce a distinct, non-tradable final good. Following Caliendo and Parro (2015), the production function is modeled as a standard CES aggregator:

$$Y_n^{k,f} = \left[\int_0^1 y_n^{k,f}(\omega^k)^{\frac{\eta-1}{\eta}} d\omega^k \right]^{\frac{\eta}{\eta-1}}, \quad \eta > 1.$$

The unit cost of producing the final good for firm f is the exact CES price index of the sourced intermediate varieties:

$$\mathbb{C}_n^k = \left[\int_0^1 \min_{j \in \mathcal{N}} \{ \mathbb{C}_{jn}^k(\omega^k) \}^{1-\eta} d\omega^k \right]^{\frac{1}{1-\eta}}, \quad (8)$$

where $\min_{j \in \mathcal{N}} \{ \mathbb{C}_{jn}^k(\omega^k) \}$ denotes the lowest available import price of intermediate good ω^k encountered by the firm across all potential origin cities.

A central departure from standard trade models is that buyers do not costlessly observe the full set of supplier price offers across space. In practice, evaluating potential suppliers, verifying quality, negotiating contracts, and sustaining procurement relationships require intercity business travel, so passenger frictions raise the effective import prices that final goods firms can achieve. We microfound this procurement friction with a costly search technology: by directing search toward city j , a buyer improves the distribution of import price offers available from that origin. Following Kortum (1997) and Eaton and Kortum (1999), we represent the underlying supplier pool with a Poisson point process over productivities; given local factor costs and iceberg freight, a better productivity draw is simply a lower import price.³ Specifically, for any variety ω^k , the number of potential intermediate suppliers in city j with productivity higher than x follows a Poisson distribution with expectation $\bar{\lambda}_j^k(x) = T_j^k x^{-\theta^P}$:

$$\Pr \{ N_j^k(x) = t \} = \frac{\bar{\lambda}_j^k(x)^t}{t!} e^{-\bar{\lambda}_j^k(x)},$$

where the scale parameter T_j^k governs the expected count of potential suppliers and the shape

³Following the original setup in Kortum (1997) and Eaton and Kortum (1999), given a time point, the total number of ideas (or potential suppliers) arriving in a location capable of producing a given variety follows a Poisson distribution. The production efficiency of each idea is independently drawn from a Pareto distribution. Mathematically, this setup is equivalent to assuming that the productivity draws are generated by a Poisson point process. We show this in detail in Appendix A.1.

parameter θ^P inversely governs the dispersion of productivity draws.

Final goods firms do not possess perfect global information to fully access this entire distribution. To source intermediates from city j , a final goods firm f in city n directs search efforts to randomly encounter a fraction $m_{nj}^k \in (0, 1)$ of these potential suppliers.⁴ By the standard thinning property of Poisson point processes, the number of encountered suppliers with productivity strictly greater than x , denoted by $N_{jn}^k(x)$, is also Poisson distributed, but with a proportionally scaled expectation $m_{nj}^k \bar{\lambda}_j^k(x)$ (see Appendix A.2 for a proof):

$$\Pr \left\{ N_{jn}^k(x) = t \right\} = \frac{\left(m_{nj}^k \bar{\lambda}_j^k(x) \right)^t}{t!} e^{-m_{nj}^k \bar{\lambda}_j^k(x)}.$$

Once the search process in city j is completed, the final goods firm observes the prices of all encountered suppliers and selects the lowest import price offer among the suppliers it has encountered in city j . Because the local factor cost c_j^k and freight cost are constant across firms in city j , selecting the lowest import price is equivalent to selecting the supplier with the highest productivity draw. Let $z_{jn}^k(\omega^k)$ denote this highest productivity among the encountered intermediate firms in city j . Its cumulative distribution function is simply the probability that exactly zero encountered firms have a productivity greater than x :⁵

$$G_{jn}^k(x) \equiv \Pr \left\{ z_{jn}^k(\omega^k) \leq x \right\} = \Pr \left\{ N_{jn}^k(x) = 0 \right\} = e^{-T_j^k m_{nj}^k x^{-\theta^P}}. \quad (9)$$

This distribution function summarizes the encountered technological frontier.

Note that when search frictions are assumed away (i.e., $m_{nj}^k = 1$ for all $n, j \in \mathcal{N}$), this distribution perfectly reduces to the classical assumption of a Fréchet distribution. In an information-frictionless economy, perfect global information ensures that buyers always match with the highest-productivity supplier in each city. In contrast, our framework implies that a relatively low-productivity firm may still secure orders if it happens to be the best match within the limited encountered set of a buyer. Compared to the classical benchmark, the scale parameter in our search-frictional economy is downscaled to $T_j^k m_{nj}^k$. Therefore, this micro-level search friction manifests as an endogenous leftward shift of the encountered technological frontier (or equivalently, an endogenous deterioration in the distribution of import price offers available from city j).

Intermediate goods markets are perfectly competitive. When shipping an intermediate good from

⁴Strictly speaking, the search effort is an individual firm's choice and should be denoted with a firm-specific superscript. However, because all final goods firms within the same sector and city are symmetric, facing identical intermediate input pools, fixed entry costs, and CES demand elasticities, they will choose the exact same optimal search effort in equilibrium. For notational convenience, we suppress the firm superscript f for the search effort and the resulting productivity distributions.

⁵To clearly see the economic intuition of this search friction, note that the probability in Equation (9) can be equivalently derived as follows without directly invoking the thinning property:

$$\Pr \left\{ z_{jn}^k(\omega^k) \leq x \right\} = \sum_{t=0}^{\infty} \Pr \left\{ N_{jn}^k(x) = t \right\} \left(1 - m_{nj}^k \right)^t = e^{-T_j^k m_{nj}^k x^{-\theta^P}}.$$

This formulation explicitly illustrates how spatial frictions induce resource misallocation: a final goods firm might source from a low-productivity supplier simply because it “missed” the $1, 2, \dots, t$ strictly better suppliers.

j to n , the supplier charges a price that covers both the marginal production cost and the standard iceberg freight transport cost $\bar{d}_{jn}$. In each city j , the import price offered by the encountered, most-productive supplier is thus $c_{jn}^k(\omega^k) = c_j^k \bar{d}_{jn} / z_{jn}^k(\omega^k)$. Then the final goods firm selects the minimum price across all cities. Given the distribution of $\{z_{jn}^k(\omega^k)\}_{j \in \mathcal{N}}$, the probability that city n imports variety ω^k from city j is:

$$\pi_{jn}^k \equiv \Pr \left\{ \frac{c_j^k \bar{d}_{jn}}{z_{jn}^k(\omega^k)} \leq \min_{j' \in \mathcal{N}} \left\{ \frac{c_{j'}^k \bar{d}_{j'n}}{z_{j'n}^k(\omega^k)} \right\} \right\} = \frac{T_j^k m_{nj}^k (c_j^k \bar{d}_{jn})^{-\theta^P}}{\gamma_n^k}, \quad (10)$$

where

$$\gamma_n^k \equiv \sum_{j'=1}^N T_{j'}^k m_{nj'}^k (c_{j'}^k \bar{d}_{j'n})^{-\theta^P}. \quad (11)$$

This probability perfectly coincides with the expenditure share (i.e., the trade share) of city n on intermediate goods sourced from city j (Eaton and Kortum, 2002). Furthermore, applying the Law of Large Numbers to Equation (8), the unit cost of the final good is given by

$$c_n^k = \left(\mathbb{E} \left[\min_{j \in \mathcal{N}} \{c_{jn}^k(\omega^k)\}^{1-\eta} \right] \right)^{\frac{1}{1-\eta}} = \gamma \left(\gamma_n^k \right)^{-\frac{1}{\theta^P}}, \quad (12)$$

where $\gamma \equiv [\Gamma((\theta^P + 1 - \eta)/\theta^P)]^{1/(1-\eta)}$ is a constant.⁶ This equation implies that the unit cost is decreasing in the search efforts $\{m_{nj}^k\}_{j \in \mathcal{N}}$. This formalizes the mechanism wherein final goods firms reduce unit costs by intensifying search toward origins that improve their import-price draws.

2.2.3 Endogenous search

Final goods firms make three interconnected decisions: they set prices under monopolistic competition, choose their optimal search efforts, and decide whether to enter the market.

First, combining Equations (6) and (7) yields the total aggregate local demand for final goods produced by firm f in sector k of city n :

$$Y_n^{k,f} = C_n^{k,f} + Q_n^{k,f} = \frac{(p_n^{k,f})^{-\sigma}}{(p_n^k)^{1-\sigma}} w_n \tilde{L}_n^k, \quad (13)$$

where $\tilde{L}_n^k \equiv \rho_n^k L_n + \sum_{s=1}^K (\gamma_n^{s,k} / \gamma_n^s) L_n^s$. Given the unit cost of the final good c_n^k , final goods firm f sets the price to maximize its gross profit:

$$\max_{p_n^{k,f}} \Pi_n^{k,f} = Y_n^{k,f} (p_n^{k,f} - c_n^k).$$

In this standard monopolistic competition setting, each individual firm sets the optimal price as a constant markup over the unit cost: $p_n^{k,f} = [\sigma/(\sigma - 1)] c_n^k$. Accordingly, the maximized gross profit

⁶Notice that this requires $\theta^P + 1 > \eta$. However, because our subsequent analysis does not explicitly depend on the value of η , we do not directly test this condition. Since this parametric restriction is standard and required for the widely adopted Eaton and Kortum (2002) framework to be well-defined, we assume it holds here.

is given by

$$\Pi_n^{k,f*} = \left(\frac{\sigma}{\sigma-1} \right)^{-\sigma} \frac{1}{\sigma-1} \frac{w_n \tilde{L}_n^k}{(p_n^k)^{1-\sigma}} (c_n^k)^{1-\sigma}. \quad (14)$$

Because $\sigma > 1$, a reduction in the unit production cost mechanically translates into a higher gross profit. This establishes the fundamental incentive for final goods firms to undertake costly travel in order to secure better import terms on intermediate inputs.

Second, firms optimally determine their search efforts $\{m_{nj}^k\}_{j \in \mathcal{N}}$. Suppose that to attain a search intensity m_{nj}^k in destination j , a firm must employ $l_{nj}^{k,f} = \varkappa (m_{nj}^k)^\kappa$ units of local labor, where $\varkappa > 0$ and $\kappa > 1$ reflect the scale and convexity of search costs. Crucially, raising search intensity in city j requires intercity business travel by the buyer, so labor costs are scaled by an iceberg passenger friction τ_{nj} . The associated cost is therefore $\tau_{nj} w_n l_{nj}^{k,f} = \varkappa \tau_{nj} w_n (m_{nj}^k)^\kappa$. In the model, τ_{nj} raises the cost of travel-mediated procurement effort that improves the import prices available from j ; search intensity m_{nj}^k is the endogenous object through which these efforts translate into better procurement terms. We do not separately model each procurement activity leading to lower achieved prices (negotiation, relationship management, etc.); instead, their combined effect on sourcing costs is summarized by the endogenous search intensity m_{nj}^k . The firm's problem is to choose $\{m_{nj}^k\}_{j \in \mathcal{N}}$ to maximize its gross profit minus total search costs:

$$\max_{\{m_{nj}^k\}_{j \in \mathcal{N}}} \Pi_n^{k,f*} - \sum_{j=1}^N \varkappa \tau_{nj} w_n (m_{nj}^k)^\kappa.$$

This optimization formalizes the core trade-off between the procurement-cost gains from better import-price draws and the travel-time costs of business travel.⁷ The first-order condition implies that the optimal search effort satisfies

$$m_{nj}^k = \left[\frac{\sigma^{-\sigma} \gamma^{1-\sigma}}{(\sigma-1)^{-\sigma} \theta^P \varkappa \kappa} \cdot \frac{\tilde{L}_n^k}{(p_n^k)^{1-\sigma}} \cdot (\gamma_n^k)^{\frac{\sigma-1-\theta^P}{\theta^P}} \cdot T_j^k (c_j^k \bar{d}_{jn})^{-\theta^P} \cdot \tau_{nj}^{-1} \right]^{\frac{1}{\kappa-1}}, \quad \forall j \in \mathcal{N}. \quad (15)$$

Notice that this equation does not yield a closed-form solution for the optimal search effort (since γ_n^k is a function of the entire search effort vector), thus we provide the discussion about its uniqueness and global optimality in Appendix A.3. Yet, substituting Equation (15) back into Equation (10) can yield the closed-form trade share under optimal endogenous search:

$$\pi_{jn}^k = \frac{\left(T_j^k\right)^{\frac{\kappa}{\kappa-1}} \left(c_j^k \bar{d}_{jn}\right)^{-\frac{\kappa \theta^P}{\kappa-1}} \tau_{nj}^{\frac{1}{1-\kappa}}}{\sum_{j'=1}^N \left(T_{j'}^k\right)^{\frac{\kappa}{\kappa-1}} \left(c_{j'}^k \bar{d}_{j'n}\right)^{-\frac{\kappa \theta^P}{\kappa-1}} \tau_{nj'}^{\frac{1}{1-\kappa}}}. \quad (16)$$

⁷To rigorously ensure a strict interior solution $m_{nj}^k \in (0, 1)$ across all bilateral pairs, we assume the search cost scale parameter $\varkappa$ is sufficiently large. Because $\varkappa$ operates as a time-invariant multiplier across the optimal search conditions, its absolute level cancels out when we focus on the relative changes of variables in the Exact Hat Algebra framework. Therefore, this theoretical bounding assumption does not impose any empirical restrictions on our subsequent counterfactual analysis.

Equation (16) highlights a fundamental mechanism of our framework: the structural decoupling and interactive amplification of physical freight costs ($\bar{d}_{jn}$) and informational search frictions (τ_{nj}). Compared to the classical Eaton and Kortum (2002) benchmark, the elasticity of trade with respect to physical freight costs is structurally amplified from θ^P to $\kappa\theta^P/(\kappa - 1) > \theta^P$. This is because, anticipating that high physical shipping costs will erode the ex-post profitability of sourcing from city j , final goods firms in city n rationally decrease their ex-ante search efforts. This reduced search systematically shrinks the effective technology pool encountered, compounding the initial physical trade barrier with an endogenous informational one.

More importantly, our framework formally separates the business travel friction τ_{nj} from the physical freight cost $\bar{d}_{jn}$. The business travel friction enters the trade share distinctively, governed by the convexity of search costs $1/(\kappa - 1)$. When HSR connects two cities, it substantially compresses business travel time without directly altering physical freight costs. This shock lowers the marginal cost of travel-mediated procurement effort, leading buyers to travel more intensively toward connected origins, improve the import prices they achieve there, and expand trade along the route even when physical freight costs $\bar{d}_{jn}$ are unchanged. HSR can therefore raise bilateral trade purely by compressing passenger frictions in the sourcing process.

2.2.4 Firm entry and equilibrium prices

To close the firm block, we endogenize the number of final goods firms F_n^k . Suppose that final goods firms can enter sector k in city n by paying a fixed entry cost, denominated in $\bar{c}_n^k$ units of local labor. In a free-entry equilibrium, the optimized net profits of each active firm are entirely dissipated by this fixed entry cost.

Because all final goods firms in the same sector of the same city are symmetric, they exert identical optimal search efforts, face the same unit cost c_n^k , and set an identical price $p_n^{k,f} = [\sigma/(\sigma - 1)]c_n^k$. According to Equation (2), the sector-specific aggregate price index simplifies to:

$$p_n^k = \frac{\sigma}{\sigma - 1} c_n^k \left(F_n^k \right)^{\frac{1}{1-\sigma}}. \quad (17)$$

Substituting this symmetric pricing rule back into the gross profit function, we obtain the maximized gross profit for each individual firm:

$$\Pi_n^{k,f*} = \frac{1}{\sigma} \frac{w_n \tilde{L}_n^k}{F_n^k}. \quad (18)$$

Additionally, utilizing the first-order condition with respect to the optimal search effort, the total optimized search costs incurred by a firm can be shown to be (see Appendix A.4 for a detailed derivation):

$$\sum_{j=1}^N \kappa \tau_{nj} w_n \left(m_{nj}^k \right)^\kappa = \frac{\sigma - 1}{\kappa \sigma \theta^P} \frac{w_n \tilde{L}_n^k}{F_n^k}. \quad (19)$$

Combining Equation (18) and Equation (19), and equating the remaining net profits to the fixed

entry cost $w_n \bar{c}_n^k$, we obtain the equilibrium number of firms:⁸

$$\Pi_n^{k,f*} - \sum_{j=1}^N \kappa \tau_{nj} w_n (m_{nj}^k)^\kappa = w_n \bar{c}_n^k \implies F_n^k = \frac{\kappa \theta^P - \sigma + 1}{\kappa \sigma \theta^P} \frac{\tilde{L}_n^k}{\bar{c}_n^k}. \quad (20)$$

This expression indicates that a larger local market size ($\tilde{L}_n^k$) or a lower entry barrier ($\bar{c}_n^k$) sustains a greater variety of final goods firms.

Moreover, substituting Equation (20) into Equation (17), we obtain the final good price index in equilibrium:

$$p_n^k = \frac{\sigma \gamma}{\sigma - 1} \left(\frac{\kappa \theta^P - \sigma + 1}{\kappa \sigma \theta^P} \right)^{\frac{1}{1-\sigma}} (\bar{c}_n^k)^{\frac{1}{\sigma-1}} (\tilde{L}_n^k)^{\frac{1}{1-\sigma}} (\gamma_n^k)^{-1/\theta^P}. \quad (21)$$

This equation captures a canonical spatial agglomeration force. As more final goods firms enter the market driven by a larger market size $\tilde{L}_n^k$, the endogenous expansion in variety drives down the final goods price index (since $1/(1 - \sigma) < 0$). This “love-of-variety” effect, which resonates with the insights emphasized by Handbury and Weinstein (2015), increases the real income derived from final goods consumption and attracts labor migration. In general equilibrium, this agglomeration force will be counterbalanced by a dispersion force from the housing market, where population inflows mechanically increase local housing rents.

Finally, by substituting Equations (15) and (21) into Equation (11), we solve for γ_n^k in equilibrium:

$$\gamma_n^k = \Lambda (\bar{c}_n^k)^{\frac{1}{\kappa}} \left[\sum_{j=1}^N (T_j^k)^{\frac{\kappa}{\kappa-1}} (c_j^k \bar{d}_{jn})^{-\frac{\kappa \theta^P}{\kappa-1}} \tau_{nj}^{\frac{1}{1-\kappa}} \right]^{\frac{\kappa-1}{\kappa}},$$

where $\Lambda \equiv (\sigma - 1)^{1/\kappa} [\kappa(\theta^P - \sigma + 1)]^{-1/\kappa}$ is a constant.

2.3 Market clearing

To close the model, we specify the market clearing conditions for housing, goods, and labor. The interplay of these markets determines the general equilibrium of the spatial economy.

2.3.1 Housing market

Following the standard approach in the urban literature, we assume that housing H_n is supplied by a competitive construction sector that uses a fixed area of land M_n and capital K_n as inputs (Epple et al., 2010; Ahlfeldt et al., 2015; Combes et al., 2019). To simplify the analysis, we assume that the production function is Cobb–Douglas:

$$H_n = K_n^\zeta M_n^{1-\zeta}, \quad 0 < \zeta < 1,$$

where $1 - \zeta$ represents the cost share of land in construction. Let $\bar{r}$ denote the exogenous unit price of capital, and p_n^{land} denote the endogenous local land rent. Cost minimization by construction firms

⁸For the equilibrium number of firms to be strictly positive, we impose the parametric restriction $\kappa \theta^P > \sigma - 1$. Our subsequent parameter estimation will verify this quantitative assumption.

yields the housing rent $r_n = \zeta^{-\zeta} (1 - \zeta)^{\zeta-1} \bar{r}^\zeta (p_n^{\text{land}})^{1-\zeta}$. Given total housing demand H_n , this housing rent can be rewritten as

$$r_n = \frac{1}{\zeta} \bar{r} M_n^{\frac{\zeta-1}{\zeta}} H_n^{\frac{1-\zeta}{\zeta}}. \quad (22)$$

On the demand side, recall from the worker's budget constraint that a fraction $1 - \beta$ of wage income is allocated to housing. This housing expenditure is collected by local landlords and further recycled into local final goods consumption. Consequently, the housing market clearing condition is given by

$$r_n H_n = (1 - \beta) w_n L_n. \quad (23)$$

Combining Equations (22) and (23), we solve for the equilibrium housing rent as a function of local population and nominal wage:

$$r_n = \zeta^{-\zeta} \bar{r}^\zeta (1 - \beta)^{1-\zeta} M_n^{\zeta-1} w_n^{1-\zeta} L_n^{1-\zeta}. \quad (24)$$

This equation captures the core dispersion force of our spatial equilibrium. Because $1 - \zeta > 0$, the inflow of population bids up local housing rents due to the constrained supply of land.

2.3.2 Goods and labor markets

Let X_j^k denote the total expenditure on final goods in sector k of city j . This expenditure originates from two sources: final consumption and input demand from all local intermediate producers. The goods market clearing condition is thus given by

$$X_j^k = \rho_j^k w_j L_j + \sum_{s=1}^K \gamma_j^{s,k} R_j^s, \quad (25)$$

where R_j^s is the total revenue of intermediate goods produced in sector s of city j .

Since final goods firms in sector s of city n set prices with a markup $\sigma/(\sigma - 1)$, their total expenditure on intermediate varieties exactly equals a fraction $(\sigma - 1)/\sigma$ of their total revenue X_n^s . Consequently, the total revenue of intermediate firms in sector s of city j is the sum of intermediate expenditures across all cities buying from them:

$$R_j^s = \sum_{n=1}^N \frac{\sigma - 1}{\sigma} \pi_{jn}^s X_n^s. \quad (26)$$

Finally, we consider the labor market clearing condition. The total local labor income $w_n L_n$ is composed of payments from two distinct sources: intermediate goods firms and final goods firms. Intermediate producers in sector k allocate a fraction γ_n^k of their total revenue to production labor. Driven by the free-entry condition, final goods firms in each sector exhaust their aggregate gross profits entirely by hiring local labor for business search and firm entry, which mathematically

amounts to $(1/\sigma)X_n^k$.⁹ Summing these two components, the labor market clears as

$$w_n L_n = \sum_{k=1}^K \left(\gamma_n^k R_n^k + \frac{1}{\sigma} X_n^k \right). \quad (27)$$

2.4 HSR shocks and the counterfactual equilibrium in relative changes

To quantify the general equilibrium impacts of the HSR network, we employ the Exact Hat Algebra (EHA) approach pioneered by Dekle et al. (2007). For any variable x , let x' denote its counterfactual value and define its relative change as $\hat{x} \equiv x'/x$. Focusing on the relative changes of endogenous variables allows us to compute counterfactual outcomes using only the observed baseline data and the relative changes in exogenous shocks.

Within this EHA framework, we introduce the HSR network as exogenous shocks that simultaneously reduce both leisure travel frictions μ_{nm} and business travel frictions τ_{nj} . Formally, we model these spatial frictions as constant-elasticity functions of actual travel time, given by $\mu_{nm} = \bar{a}^T t_{nm}^{\alpha^T}$ and $\tau_{nj} = \bar{a}^P t_{nj}^{\alpha^P}$. We allow the elasticity parameters (α^T and α^P) to differ, reflecting the distinct time sensitivities of residents engaging in leisure activities versus final goods firms conducting business travel for procurement. Consequently, a reduction in travel time translates into relative changes in spatial frictions, given by $\hat{\mu}_{nm} = \hat{t}_{nm}^{\alpha^T}$ and $\hat{\tau}_{nj} = \hat{t}_{nj}^{\alpha^P}$.

By quantifying the actual travel time reductions induced by HSR, we can directly feed these exogenous travel time shocks $\{\hat{t}_{nj}\}_{n,j \in \mathcal{N}}$ into the system to solve for the full counterfactual spatial equilibrium in relative changes. Because the EHA method has evolved into a standard method in the literature, we provide the formal definition of the equilibrium in relative changes and the full system of non-linear equations in Appendix A.5.

3 Data and HSR network evolution

3.1 Overview of the data construction

To calibrate the structural model and implement the counterfactual analysis, we compile a comprehensive set of datasets covering over 300 cities in China from 2008 to 2015.¹⁰ In the main text, we provide an overview of our primary data sources and the construction of key variables. Detailed data processing procedures and justifications are provided in Appendix B.

⁹This can be formally shown as follows. First, Equation (18) indicates that the gross profit of all final goods in sector k of city n is given by $(1/\sigma)w_n \bar{L}_n^k$. Second, multiplying both sides of Equation (13) by $p_n^{k,f}$ and summing over all final goods firms f yields the total sector revenue $X_n^k = w_n \bar{L}_n^k$. Therefore, the aggregate gross profit in sector k of city n is exactly $(1/\sigma)X_n^k$, which is fully transferred to workers as payments for search and entry costs.

¹⁰China's administrative hierarchy comprises 34 provincial-level regions, including 4 directly-administered municipalities (provincial-level cities): Beijing, Tianjin, Shanghai, and Chongqing. The remaining provinces are further subdivided into 333 prefecture-level administrative regions, consisting of 293 prefecture-level cities, 7 prefectures, 30 autonomous prefectures, and 3 leagues. Throughout our analysis, the term "city" refers to either a direct-administered municipality or a prefecture-level administrative region (thus 337 cities in total). Because our counterfactual simulations require the integration of all datasets, we restrict the spatial coverage in the counterfactual analysis to the intersection of all available locations, yielding a set of 302 cities.

Demographic data.—To determine the center of each city, we utilize gridded demographic data from the Gridded Population of the World (GPWv4) dataset, specifically Version 4.11, provided by NASA. This dataset provides population count estimates per grid cell, based on the 2010 National Population Census in China, at a spatial resolution of $0.0083^\circ \times 0.0083^\circ$. We use the population count of each cell as a weight to calculate the “center of mass” for each city’s polygon, which we subsequently designate as the city population center.

Highway network data.—To measure the evolution of road connectivity, we utilize geographic data on China’s highway and first-rate road networks spanning 2008 to 2015, originally constructed by Ma and Tang (2024).¹¹ The raw data provide a digitized, pixel-level mapping of the road network along with pixel-specific engineering design speeds ranging from 60 km/h to 120 km/h. To calculate city-to-city pairwise travel times, we overlay this network onto a high-resolution $0.005^\circ \times 0.005^\circ$ spatial grid and employ a grid-based least-cost path algorithm to identify the time-minimizing route between any two city centers. This baseline highway travel time serves as the benchmark to quantify the travel time reductions induced by the HSR network. Furthermore, to construct a corresponding measure for road freight travel times, we impose a speed limit of 100 km/h and recompute the minimum pairwise travel times.¹²

Railway schedule snapshots.—Our primary measure of railway accessibility relies on snapshots of official train schedules for eight distinct dates spanning 2008 to 2015. The raw data provide sequential timetables for all operational passenger trains, detailing the ordered sequence of stops alongside precise arrival and departure times. Because the Chinese intercity railway system operates under a rigid, centrally planned scheduling framework, a single-day snapshot reliably reflects the stable transport connectivity during that period. We map over 4,000 unique railway stations to their corresponding prefecture-level cities and geographic coordinates. We then convert these sequential timetables into a bilateral station-to-station origin–destination dataset. Following standard empirical practice, we utilize the prefix of each train number to distinguish HSR (prefixes G, D, and C) from conventional trains. Finally, to ensure strict comparability across transportation modes, we construct a measure of effective door-to-door travel time by combining the in-vehicle rail transit time with the local road travel time from the respective city centers to the stations. The city-to-city pairwise railway travel time—calculated separately for HSR and conventional trains—is defined as the minimum door-to-door travel time across all available schedules directly connecting any given city pair.

Migration data.—To construct the bilateral labor migration matrix and measure spatial labor mobility, we utilize individual-level microdata from the 2010 National Population Census and the 2015 National 1% Population Sampling Survey. These datasets consistently report respondents’ current residential location and their official household registration (*hukou*) location.¹³ Consistent with many quantitative studies about China (e.g., Tombe and Zhu, 2019; Liang et al., 2024), we

¹¹We classify both highways and first-rate roads under the unified category of the highway network. For expositional simplicity, we refer to them collectively as “highways” throughout the text.

¹²According to the *Implementing Regulations of the Road Traffic Safety Law of China*, large motor vehicles (including freight trucks) are strictly prohibited from exceeding 100 km/h on highways. The document is available at https://www.gov.cn/zwqk/2005-05/23/content_214.htm (accessed July 2026).

¹³For a more detailed description of the *hukou* system in China, see Song (2014).

define an individual’s origin city by their *hukou* location and their destination city by their current residence. Migrants are then defined as individuals whose *hukou* location is different from their current residence. By aggregating these micro-level records to the prefecture-city pair level, we construct the directional migration shares.

Intercity, multi-sector input-output data.—To capture spatial production networks and calibrate the structural elasticity parameters, we utilize the China City-Level Multi-Regional Input-Output (MRIO) tables for the years 2012 and 2015, provided by Zheng et al. (2022). The dataset estimates the intermediate input flows, value added, and final demands across 42 economic sectors among 313 regions. We leverage these matrices to directly construct the sector-specific bilateral trade shares across cities and calibrate consumption shares and production elasticities.

UnionPay data.—We measure bilateral business and leisure travel using nationwide, anonymized credit- and debit-card transactions processed through the UnionPay network from 2011 to 2015. During this period, UnionPay was the de facto monopoly provider of domestic interbank card-clearing services, and its underlying records therefore covered virtually all interbank card transactions in China. Our dataset contains complete transaction histories for a random 1% sample comprising approximately 27 million cards, providing unusually broad longitudinal coverage of intercity passenger movements. We retain offline transactions, which record purchases made while the cardholder is physically present at the merchant, and map merchant locations to prefecture-level cities. We assign each card a home city based on its most frequent location of use within a rolling twelve-month window and identify travel when the card is used outside that city. We classify trips on non-holiday weekdays as business travel and trips during national public holidays and an extended Chinese New Year window as leisure travel.

3.2 HSR network evolution

Before proceeding to the structural estimation, we document two stylized facts regarding the spatial expansion and long-haul economy of China’s HSR network. These facts not only illustrate the real-world infrastructure planning process but also provide empirical foundations for our subsequent counterfactual simulations.

Fact #1: Volume-oriented network expansion. The deployment of the HSR network was not geographically random; instead, it exhibited a notable volume-oriented prioritization. To illustrate this, we utilize the 2011 UnionPay bilateral business travel volume as a time-invariant baseline measure of intercity passenger demand. As shown in Table 1, during the initial phase of the network in 2008, the system connected only 293 directed city pairs, yet these pairs accounted for 16% of the total national business travel volume. As the network aggressively expanded over the years, the mean distance between connected cities gradually increased. We graphically illustrate this connection evolution in Appendix Figure E.1, which demonstrates that the network initially focused on the highly developed eastern regions before gradually penetrating into the central and western regions. By 2015, the network had grown to cover 6,669 pairs, accompanied by a significant decline in the average volume per connected pair. This declining average suggests that as the network extended into less dense regions, marginal infrastructure investments likely faced an efficiency

trade-off.

This volume-oriented expansion is further reflected in the network’s coverage of the highest-demand corridors. In 2008, the network covered 26 of the top 0.1% (90 pairs in total) and 118 of the top 1% (about 900 pairs in total) highest-volume routes. By 2015, these numbers had reached 82 and 601, respectively. However, the fact that nearly one-third of the top 1% routes remained unconnected by 2015 implies that passenger volume, while crucial, was not the exclusive determinant of network expansion. A growing body of political economy literature highlights that China’s HSR rollout was also deeply shaped by central–local bargaining and local officials’ promotion incentives (e.g., Ma, 2022). These strategic political maneuvers may have led to deviations from an efficiency-maximizing network. We will contrast the actual network expansion with a strictly volume-based construction order in our counterfactual simulations.

Table 1: The evolution of China’s HSR network from 2008 to 2015

Year	# connected pairs	Mean distance (km)	Total volume share (%)	Volume per pair	# top 0.1% pairs	# top 1% pairs
2008	293	343.27	16.19	2,285.13	26	118
2009	1,073	430.18	23.64	910.91	31	221
2010	1,461	409.86	31.45	890.09	39	310
2011	1,735	426.36	37.88	902.64	48	342
2012	2,461	474.49	46.04	773.54	56	401
2013	2,960	510.65	49.48	691.24	59	415
2014	5,231	643.35	60.37	477.17	74	525
2015	6,669	667.91	66.74	413.81	82	601

Notes: This table summarizes the evolution of China’s HSR network from 2008 to 2015. The unit of observation is a directed city pair ($337 \times 336 = 113,232$ in total), which accounts for rare instances where outbound and inbound services have asymmetric intermediate stops. Columns (2) and (3) report the total number of connected pairs and the mean distance between their city centers, respectively. To measure intercity passenger demand, we utilize 2011 UnionPay business travel volume as a time-invariant baseline measure. Based on this baseline, Columns (4) and (5) report the share of total travel volume captured by the HSR-connected pairs and the average volume per pair. Columns (6) and (7) present the number of connected pairs that rank within the top 0.1% and 1% of all directed routes in terms of baseline travel volume.

Fact #2: The long-haul economy of HSR. According to Koster et al. (2022), a transportation technology exhibits a “long-haul economy” if travel time (or cost) increases less than proportionately with geographical distance. To empirically quantify this feature across different transportation modes, we estimate the following specification:

$$\ln t_{njt} = \alpha_t + \beta_t \ln \text{Distance}_{nj} + \epsilon_{njt}, \quad (28)$$

where t_{njt} is the minimum travel time between city n and city j in year t by a given transportation mode, α_t is the year-specific intercept, Distance_{nj} is the geographic distance between the centers of city n and city j , β_t is the distance elasticity in year t , and ϵ_{njt} is the error term.

Table 2 reports the estimation results for HSR, conventional trains, and highways. Consistent with the findings in Koster et al. (2022), the distance elasticity of HSR travel time (β_t) is consistently and significantly less than one across all observed years (e.g., 0.781 in 2015). Economically, this implies that a 1% increase in physical distance translates to only a 0.781% increase in HSR travel

time. In contrast, the elasticities for conventional trains and highways are much closer to unity (averaging approximately 0.91 and 0.97, respectively).

This structural advantage over long distances is a defining feature of the HSR network. Accurately documenting this relationship is critical for our counterfactual analysis. Rather than assuming a naive, uniform proportional time reduction across all city pairs, incorporating this empirically estimated long-haul economy allows us to realistically parameterize and predict travel times for unobserved routes when simulating a fully connected national network.

Table 2: The long-haul economy of HSR, conventional train, and highway

	Dependent variable: Logarithm of travel time		
	High-speed rail	Conventional train	Highway
	(1)	(2)	(3)
$\ln(\text{distance}) \times \mathbb{1}\{t = 2008\}$	0.678*** (0.025)	0.892*** (0.008)	0.963*** (0.009)
$\ln(\text{distance}) \times \mathbb{1}\{t = 2009\}$	0.769*** (0.018)	0.913*** (0.008)	0.968*** (0.009)
$\ln(\text{distance}) \times \mathbb{1}\{t = 2010\}$	0.791*** (0.021)	0.910*** (0.007)	0.970*** (0.009)
$\ln(\text{distance}) \times \mathbb{1}\{t = 2011\}$	0.755*** (0.019)	0.911*** (0.007)	0.970*** (0.009)
$\ln(\text{distance}) \times \mathbb{1}\{t = 2012\}$	0.755*** (0.019)	0.910*** (0.007)	0.965*** (0.008)
$\ln(\text{distance}) \times \mathbb{1}\{t = 2013\}$	0.711*** (0.016)	0.925*** (0.007)	0.967*** (0.008)
$\ln(\text{distance}) \times \mathbb{1}\{t = 2014\}$	0.767*** (0.014)	0.920*** (0.007)	0.967*** (0.008)
$\ln(\text{distance}) \times \mathbb{1}\{t = 2015\}$	0.781*** (0.012)	0.914*** (0.007)	0.970*** (0.008)
Year fixed effects	Yes	Yes	Yes
# of observations	21,883	191,551	905,856

Notes: This table reports the estimation results for Equation (28). The estimation sample comprises a panel of bilateral travel times spanning the period 2008 to 2015; consequently, the coefficients on the interaction terms capture the year-specific elasticities of travel time with respect to distance. Standard errors, two-way clustered at the origin and destination city levels, are reported in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

4 Parameter calibration and estimation

As shown in our definition of the counterfactual equilibrium, implementing the counterfactual analysis requires parameterizing the model. The set of required structural parameters includes $\{\alpha^T \theta^T, \zeta, \beta, \theta^R, \psi / \theta^T, \theta^P, \kappa, \alpha^P, \sigma\}$ and a set of Cobb–Douglas shares $\{\rho_n^k, \gamma_n^k, \gamma_n^{k,s}\}_{n \in \mathcal{N}, k, s \in \mathcal{K}}$. In this section, we first calibrate part of them based on the literature and observed data. We then estimate the remaining elasticities sequentially, exploiting the precise network variations and structural relationships outlined in the model.

4.1 Calibration and external parameters

We calibrate the consumption shares (ρ_n^k) and the production elasticities (γ_n^k for labor and $\gamma_n^{k,s}$ for materials) directly from the 2012 China City-Level MRIO table, as detailed in Appendix B.5. Moreover, we calibrate the expenditure share on final goods to $\beta = 0.85$ based on the *China Statistical Yearbook*, which represents the average share across urban and rural residents.

Next, we assign values to several structural parameters well established in the literature. First, Wu et al. (2012) estimates that land transaction prices account for, on average, 37% of total house values. Since $1 - \zeta$ quantifies the share of land costs in total construction costs (and housing rent in a competitive market), we set $\zeta = 0.63$. Second, following Ma and Tang (2024), we set the elasticity of substitution for firm-specific final goods aggregation to $\sigma = 6$, which centers within the commonly estimated range reported by Anderson and van Wincoop (2004). Third, we calibrate the shape parameter governing the spatial dispersion of residential preference shocks to $\theta^R = 1.9$, according to the intercity migration elasticity estimated by Fang and Huang (2022).

Finally, we discuss the calibration of the trade elasticity. While there is a large literature on international trade elasticities, precise estimates for intra-national trade elasticities remain scarce (Tombe and Zhu, 2019). Unlike international trade, where asymmetric bilateral tariffs provide clear observable variations and price data are widely available (e.g., the method proposed by Caliendo and Parro, 2015), precisely quantifying bilateral trade frictions between domestic cities is empirically challenging.

We benchmark our parameters against Che et al. (2024), who estimate 16 sector-specific, intra-national trade elasticities using price data in China. Under their quantitative trade model building on Caliendo and Parro (2015), the structural relationship is log-linear: $\ln(\pi_{jn}^k/\pi_{jj}^k) = -\theta^P \ln(\mathbb{C}_j^k \bar{d}_{jn}/\mathbb{C}_n^k)$. The authors then approach the right-hand side using observed spatial price gaps. However, by explicitly micro-founding the business search process, our spatial framework yields a different structural relationship:

$$\ln\left(\frac{\pi_{jn}^k}{\pi_{jj}^k}\right) = \frac{1}{\kappa - 1} \ln\left(\frac{\bar{\mathbb{C}}_n^k}{\bar{\mathbb{C}}_j^k}\right) - \frac{\kappa\theta^P}{\kappa - 1} \ln\left(\frac{\mathbb{C}_j^k \bar{d}_{jn}}{\mathbb{C}_n^k}\right) - \frac{1}{\kappa - 1} \ln\left(\frac{\tau_{nj}}{\tau_{jj}}\right).$$

This model-implied relationship indicates that, if we estimate the identical empirical specification as in Che et al. (2024), the structural interpretation of the estimated coefficient on the relative cost term maps to $-\kappa\theta^P/(\kappa - 1)$ rather than $-\theta^P$. Consequently, for our baseline calibration, we take the average of their 16 sector-specific estimates, which is 5.94, to calibrate the composite term $\kappa\theta^P/(\kappa - 1)$ as a whole. We will then solve for κ and θ^P separately in the following estimation. Furthermore, our theoretical framework suggests that estimating the identical specification inevitably relegates the unobserved relative search cost τ_{nj}/τ_{jj} into the error term. Because spatial search frictions are positively correlated with physical freight transport costs $\bar{d}_{jn}$, this omission could induce a systematic upward bias in absolute magnitude.

Reassuringly, in the counterfactual equilibrium system, this calibrated value governs how changes in suppliers' production costs $\hat{\mathbb{C}}_j^k$ reallocate trade across origins (see Appendix Equation (A.5)). By contrast, HSR enters trade shares directly through $\hat{t}_{nj}$ with elasticity $\alpha^P/(1 - \kappa)$, which does not

depend on this calibration. Therefore, we expect our results to be insensitive to a smaller value of this calibration. In our counterfactual analysis, we explicitly explore alternative, smaller values of this calibration as a robustness check.

4.2 Estimation

4.2.1 Leisure travel gravity

We start by estimating the leisure travel gravity equation to identify the composite travel friction elasticity, $\alpha^T \theta^T$. According to the structural gravity relationship implied by Equation (3), we specify the following regression equation:

$$\lambda_{m|n,t}^T = \exp \left(-\alpha^T \theta^T \ln t_{nmt} + \delta_{mt} + \delta_{nt} \right) + \epsilon_{nmt}, \quad m \in \mathcal{N} \setminus \{n\}, \quad (29)$$

where $\lambda_{m|n,t}^T$ is the share of observed intercity leisure trips from city n to city m , relative to all observed intercity leisure trips originating in n in year t , and t_{nmt} is the minimum travel time between city m and n across all transportation modes covered by our data (i.e., highway, conventional train, and HSR). The destination–time fixed effects, δ_{mt} , control for the time-varying destination attractiveness, while the origin–year fixed effects, δ_{nt} , capture the time-varying leisure travel accessibility. Finally, the observed UnionPay share may deviate from this population counterpart because it is constructed from a finite sample of transactions. We therefore add an error term ϵ_{nmt} to reflect the mismeasurement of the dependent variable and assume it is conditionally mean-zero.

To estimate the parameters, we employ the Poisson Pseudo-Maximum Likelihood (PPML) estimator. This approach naturally accommodates observed zero-flow city pairs without requiring arbitrary data transformations. Table 3 reports the estimation results. Because our UnionPay dataset spans from 2011 to 2015, we construct the bilateral leisure travel shares and travel times for each year and pool them in Column (1). The point estimate is statistically significant and economically large, suggesting that a 1% increase in bilateral travel time reduces the leisure travel share by about 2%. In columns (2)–(6), the year-specific elasticity estimates range from 1.812 to 2.111 and decline monotonically over the sample period. Nevertheless, each estimate remains close to the pooled estimate. We then set $\alpha^T \theta^T = 2$ as our baseline value.

Because our dataset omits air travel times, relying exclusively on the minimum travel time across the highway and railway networks may systematically overestimate the true travel friction for distant city pairs. Fang et al. (2025) find that HSR competition with airlines is strongest within the range of approximately 1,200 km. Motivated by this benchmark, we re-estimate the leisure travel gravity equation using restricted samples that retain only city pairs whose straight-line distances do not exceed 600, 800, 1,000, 1,200, or 1,400 km. As shown in Appendix Table E.1, the estimates remain quantitatively close to our baseline value.

Under the maintained assumption that the UnionPay destination shares are representative of the underlying distribution of intercity leisure trips, the estimated origin and destination fixed effects admit structural interpretations. Denoting the estimated fixed effects as $\tilde{\delta}_{mt}$ and $\tilde{\delta}_{nt}$, we recover each city’s destination attractiveness and leisure travel accessibility relative to the geometric averages.

Table 3: The estimation of the leisure travel gravity equation

	Dependent variable: $\lambda_{m n,t}^T$					
	All	2011	2012	2013	2014	2015
	(1)	(2)	(3)	(4)	(5)	(6)
$\ln t_{nmt}$	-1.981*** (0.089)	-2.111*** (0.068)	-2.056*** (0.075)	-2.011*** (0.094)	-1.897*** (0.112)	-1.812*** (0.106)
Destination FEs	No	Yes	Yes	Yes	Yes	Yes
Origin FEs	No	Yes	Yes	Yes	Yes	Yes
Destination $\times$ year FEs	Yes	No	No	No	No	No
Origin $\times$ year FEs	Yes	No	No	No	No	No
# of observations	559,450	111,890	111,890	111,890	111,890	111,890

Notes: This table presents the estimation results of Equation (29) using the PPML estimator. Column (1) uses the panel data of all city pairs over 2011 to 2015. Columns (2)–(6) estimate the model separately for each year. Standard errors, two-way clustered at the origin and destination city levels, are reported in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

Specifically, by centering the estimated fixed effects, we have

$$\exp\left(\tilde{\delta}_{mt} - \frac{1}{|\tilde{\mathcal{N}}_t^{\text{des}}|} \sum_{m' \in \tilde{\mathcal{N}}_t^{\text{des}}} \tilde{\delta}_{m't}\right) \rightarrow_p \left(\frac{b_{mt}}{\bar{b}_t}\right)^{\theta^T} \quad (30)$$

and

$$\exp\left[-\left(\tilde{\delta}_{nt} - \frac{1}{|\tilde{\mathcal{N}}_t^{\text{ori}}|} \sum_{n' \in \tilde{\mathcal{N}}_t^{\text{ori}}} \tilde{\delta}_{n't}\right)\right] \rightarrow_p \frac{\Phi_{nt}^T}{\bar{\Phi}_t^T}, \quad (31)$$

where $\tilde{\mathcal{N}}_t^{\text{des}}$ and $\tilde{\mathcal{N}}_t^{\text{ori}}$ are the sets of cities which have estimated destination and origin fixed effects, respectively, and $\bar{b}_t$ and $\bar{\Phi}_t^T$ denote the geometric averages of destination attractiveness and origin accessibility across all available cities in year t . We employ these recovered components, denoted by $\overline{b_{mt}^{\theta^T}}/\bar{b}_t^{\theta^T}$ and $\overline{\Phi_{nt}^T}/\bar{\Phi}_t^T$, in our subsequent estimation of the migration gravity equation. Additionally, we map the spatial distributions of relative destination attractiveness and leisure travel accessibility across cities in Appendix Figures E.2 and E.3.

Finally, using each city's observed total outbound leisure travel, we apply the estimated gravity equation to predict its total inbound volume as a destination and compare it with the observed inbound volume in Appendix C.3. The result shows a strong correlation between the predicted and observed volumes, suggesting a good fit of the gravity equation estimates in explaining the leisure travel across cities.

4.2.2 Migration gravity

Turning from leisure travel to migration, we estimate the migration gravity to identify the migration elasticity with respect to leisure travel accessibility. According to the equilibrium migration share

in Equation (5), we specify the following regression equation:

$$\ln \lambda_{jnt} - \theta^R \ln \frac{w_{nt}}{P_{nt}} = \theta^R \frac{\psi}{\theta^T} \ln \frac{\Phi_{nt}^T}{\bar{\Phi}_t^T} + \beta_t \ln (\text{Distance}_{jn} + 1) + \delta_{jt} + \epsilon_{jnt}. \quad (32)$$

Because our migration data cover 2010 and 2015, we pool the two cross-sections and introduce a time dimension into the estimating equation. In this regression, we utilize the relative accessibility term $\widehat{\Phi_{nt}^T / \bar{\Phi}_t^T}$ recovered from the previous leisure travel gravity estimation.¹⁴ Because local real wages are jointly determined with migration in equilibrium, we do not attempt to estimate their coefficient from the migration equation. Instead, we impose the calibrated $\theta^R = 1.9$ and move the real wage term to the left-hand side. This transformation eliminates the need to instrument for real wages. Moreover, we control for bilateral physical distance to capture distance-related spatial migration costs, and include origin–year fixed effects δ_{jt} to fully absorb the multilateral resistance term $\ln \Theta_{jt}$. Consequently, the structural error term ϵ_{jnt} encapsulates unobserved local amenities a_{nt} and residual migration costs.

A prominent identification challenge in this specification is that unobserved local amenities may be systematically correlated with local transport infrastructure. Because our model defines accessibility as $\Phi_{nt}^T \equiv \sum_{m' \in N \setminus \{n\}} (b_{m't} / \mu_{nm't})^{\theta^T}$, the potential correlation between a_{nt} and $\mu_{nm't}$ could induce omitted-variable bias.

To isolate variation in leisure accessibility that does not arise directly from realized bilateral transport provision, we construct an instrumental variable that replaces actual travel frictions with distance-based frictions under a uniform and time-invariant travel speed:

$$\tilde{\varphi}_{nt}^T = \sum_{m \in N \setminus \{n\}} \left(\widehat{b_{m't}^{\theta^T} / \bar{b}_t^{\theta^T}} \right) \text{Distance}_{nm}^{-\alpha^T \theta^T}, \quad (33)$$

where $\widehat{b_{m't}^{\theta^T} / \bar{b}_t^{\theta^T}}$ is the recovered relative destination attractiveness.¹⁵ Therefore, this instrument cannot increase mechanically because city n itself becomes more attractive or because it endogenously receives faster transport infrastructure. Instead, the IV captures whether city n is geographically positioned close to a more attractive portfolio of external leisure destinations. Much of the persistent component of external destination attractiveness is rooted in natural, historical, and cultural endowments outside city n 's control. These features make it plausible that the IV is orthogonal to city n 's unobserved residential amenities and residual migration costs. We thus use $\ln \tilde{\varphi}_{nt}^T$ as an instrument for $\ln (\widehat{\Phi_{nt}^T / \bar{\Phi}_t^T})$.

We report the regression results using the 2010 cross-section, the 2015 cross-section, and the

¹⁴Because the UnionPay data are only available starting in 2011, we utilize the recovered 2011 measures of leisure travel accessibility and destination attractiveness as the closest available proxies for their 2010 counterparts. For 2015, we employ the contemporaneous measures. Furthermore, because fixed effects estimation cannot identify the absolute level of Φ_{nt}^T , we rely on its relative magnitude. Conveniently, upon logarithmic transformation, the normalizing denominator $\ln \bar{\Phi}_t^T$ enters additively and is perfectly absorbed by the origin–year fixed effects.

¹⁵Specifically, assuming a uniform, time-invariant travel speed $\bar{v}$ between all city pairs, the travel time between two cities n and m is $\text{Distance}_{nm} / \bar{v}$. Therefore, we construct a counterfactual friction $\bar{\mu}_{nm} = \bar{a}^T \text{Distance}_{nm}^{\alpha^T} \bar{v}^{-\alpha^T}$ and use it to replace the actual $\mu_{nm't}$ in the definition of Φ_{nt}^T . The constants $\bar{a}^T$ and $\bar{v}$ do not provide any cross-sectional or temporal variation and therefore do not affect the IV estimation, thus we omit them from the construction of the instrument.

pooled panel in Table 4. Across all specifications, the coefficients on leisure travel accessibility are consistently positive and statistically significant. The first-stage F-statistics indicate a strong relevance between the instrument and recovered leisure accessibility. The IV estimates are modestly larger than their OLS counterparts, although the differences are quantitatively small.

For the subsequent counterfactual analysis, we adopt the estimated parameter $\theta^R \psi / \theta^T = 0.18$ from Column (6). Combined with our calibration of $\theta^R = 1.9$, this yields a value of 0.095 for ψ / θ^T . This implies that the migration elasticity with respect to nonlocal leisure accessibility is approximately 9.5% as large as the elasticity with respect to local real wages and intra-city amenities. Although this relative weight is modest, large HSR-induced changes in leisure accessibility can still make the direct leisure channel quantitatively important, as we assess in the counterfactual analysis.

Table 4: The estimation of the migration gravity equation

	Dependent variable: $\ln \lambda_{jnt} - \theta^R \ln \frac{w_{nt}}{P_{nt}}$					
	2010		2015		Both years	
	(1)	(2)	(3)	(4)	(5)	(6)
$\ln(\overline{\Phi_{nt}^T} / \overline{\Phi_t^T})$	0.101** (0.049)	0.110** (0.054)	0.209*** (0.068)	0.248*** (0.071)	0.151*** (0.055)	0.180*** (0.060)
$\ln(\text{distance})$	-0.909*** (0.015)	-0.910*** (0.016)	-0.876*** (0.016)	-0.879*** (0.017)		
$\ln(\text{distance}) \times \mathbb{1}\{t = 2010\}$					-0.914*** (0.016)	-0.916*** (0.017)
$\ln(\text{distance}) \times \mathbb{1}\{t = 2015\}$					-0.872*** (0.016)	-0.874*** (0.016)
Origin fixed effects	Yes	Yes	Yes	Yes	No	No
Origin $\times$ year fixed effects	No	No	No	No	Yes	Yes
Estimator	OLS	IV	OLS	IV	OLS	IV
First-stage F-statistic		1,759		1,039		1,594
# of observations	14,216	14,219	14,227	14,227	28,443	28,443

Notes: This table presents the estimation results of Equation (32). Columns (1) and (2) use the cross-sectional data in 2010, columns (3) and (4) use the cross-sectional data in 2015, and columns (5) and (6) use the pooled data over both years. The OLS estimator is used in Columns (1), (3), and (5), while the IV estimator is used in Columns (2), (4), and (6). The instrument is constructed according to Equation (33). Kleibergen–Paap F-statistics are reported. Standard errors, two-way clustered at the origin and destination city levels, are reported in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

We further examine the sensitivity of our estimates to the calibration of θ^R in Appendix C.1. The relevant value of the migration elasticity depends on the spatial choice margin and the definition of location alternatives. Our baseline calibration follows Fang and Huang (2022), whose city-pair migration setting and spatial resolution most closely match our analysis. The Appendix analytically characterizes the relationship between θ^R and the estimated ψ / θ^T , and re-estimates the migration gravity equation for alternative values of θ^R ranging from 1.5 to 3. As θ^R increases, the model attributes a larger share of migration responses to real-income differences, leaving a smaller estimated contribution of leisure accessibility. Although the confidence interval includes zero under the conservative upper-end calibration of $\theta^R = 3$, the point estimate remains positive throughout the

full range and is statistically significant for all reported calibrations from 1.5 through 2.5, which encompasses the values most directly comparable to our intercity migration setting.

As an additional robustness check, Appendix C.2 re-estimates the migration gravity equation in first-differences between 2010 and 2015. This specification identifies $\theta^R\psi/\theta^T$ from within-city changes in leisure travel accessibility. The IV estimate is 0.144, which is quantitatively close to our pooled baseline of 0.180, although less precisely estimated because first-differencing requiring positive migration shares in both census years substantially reduces the log-share sample and removes much of the cross-sectional variation in accessibility. Moreover, augmenting the regression with the change in bilateral travel time yields a coefficient that is small and statistically insignificant, while the accessibility estimate remains unchanged. Therefore, it provides little evidence that the home-tie mechanism drives the estimated leisure-accessibility coefficient.¹⁶

Finally, we assess the goodness of fit of the migration gravity estimates in Appendix C.3. Because Equation (32) does not estimate the complete structural migration probabilities, we reconstruct predicted migration shares from the structural definition in Equation (5), retaining observed real wages, recovered leisure accessibility, and the estimated distance elasticity, while shutting down amenities and residual non-distance migration costs. Aggregating these predicted shares with each origin's observed *hukou* population yields a predicted residential population for each destination, which we then compare with the observed population. Again, we find a strong correlation between the predicted and observed values.

4.2.3 Business search volume

We next estimate the elasticity of business travel frictions with respect to travel time, α^P . This parameter governs how reductions in passenger travel time lower the marginal cost of spatial search and thereby reshape bilateral searching and sourcing.

Recall that a final goods firm in sector k of city n must employ $l_{nj}^{k,f} = \kappa(m_{nj}^k)^\kappa$ units of local labor to attain search intensity m_{nj}^k in supplier city j , and that the corresponding search cost is $\tau_{nj}w_n l_{nj}^{k,f}$. As shown in Appendix A.4, the optimal search costs are allocated across destinations in proportion to trade shares:

$$\kappa\tau_{nj}w_n \left(m_{nj}^k\right)^\kappa = \frac{\sigma-1}{\kappa\sigma\theta^P} \frac{w_n \tilde{L}_n^k}{F_n^k} \pi_{jn}^k.$$

Therefore, the aggregate search labor traveling from buyer city n to supplier city j in sector k is

$$F_n^k l_{nj}^{k,f} = \frac{\sigma-1}{\kappa\sigma\theta^P} \pi_{jn}^k \tilde{L}_n^k \tau_{nj}^{-1}.$$

Aggregating across sectors and imposing $\tau_{nj} = \bar{a}^P t_{nj}^{\alpha^P}$ yields the model-implied total business

¹⁶A related concern is that HSR may facilitate return visits from residence city n to a migrant's origin or *hukou* city j , thereby directly reducing the effective cost of choosing the j - n migration pair. Aggregate leisure flows used to recover Φ_n^T include return-home travel and therefore capture part of its average value. However, Φ_n^T is not indexed by the migrant's origin city and need not fully absorb this origin-specific home-tie mechanism. We therefore examine the bilateral channel separately by adding the change in pairwise travel time to the first-difference specification.

search volume from n to j :

$$\sum_{k=1}^K F_n^k l_{nj}^{k,f} = \frac{\sigma - 1}{\bar{a}^P \kappa \sigma \theta^P} t_{nj}^{-\alpha^P} \sum_{k=1}^K \pi_{jn}^k \tilde{L}_n^k. \quad (34)$$

Using the definition $\tilde{L}_n^k \equiv \rho_n^k L_n + \sum_{s=1}^K (\gamma_n^{s,k} / \gamma_n^s) L_n^s$ and denoting the sectoral production-labor share by $v_n^s \equiv L_n^s / L_n$, we can rewrite

$$\tilde{L}_n^k = D_n^k L_n,$$

where $D_n^k \equiv \rho_n^k + \sum_{s=1}^K \gamma_n^{s,k} v_n^s / \gamma_n^s$.¹⁷ We then have

$$\frac{\sum_{k=1}^K F_n^k l_{nj}^{k,f}}{\sum_{k=1}^K \pi_{jn}^k D_n^k} = \frac{\sigma - 1}{\bar{a}^P \kappa \sigma \theta^P} t_{nj}^{-\alpha^P} L_n$$

Taking the logarithm and difference across years, we obtain the following regression equation:

$$\Delta \ln \frac{\sum_{k=1}^K F_n^k l_{nj}^{k,f}}{\sum_{k=1}^K \pi_{jn}^k D_n^k} = -\alpha^P \Delta \ln t_{nj} + \Delta \delta_n + \Delta \epsilon_{nj}, \quad (35)$$

where fixed effects $\Delta \delta_n \equiv \Delta \ln L_n$ and $\Delta \epsilon_{nj}$ denotes the independent error term.

However, a practical challenge remains. Although the firm block of our model can be read more broadly as travel-mediated procurement—with search providing one disciplined microfoundation for how passenger time enters sourcing costs—UnionPay business transactions do not identify trip purpose. Observed volume therefore mixes procurement-related travel with residual business trips that our sourcing block does not cover.¹⁸ For example, reverse advertising visits (suppliers in j traveling to buyers in n) and trips that primarily generate knowledge spillovers: the latter would shift the technological frontier and benefit all importers, rather than exclusively improving the specific price distribution encountered by final-good producers in buyer-city n . Moreover, residual volume also includes comparatively time-inelastic corporate travel: recurring board or shareholder meetings, pre-scheduled industry conferences, mandatory regulatory inspections, etc.

To make this problem explicit, decompose observed business travel V_{nj} into a procurement-related component that obeys the structural relationship in Equation (35) and a residual component V_{nj}^U :

$$V_{nj} = V_{nj}^U + \sum_{k=1}^K F_n^k l_{nj}^{k,f} = (1 + \phi_{nj}) \sum_{k=1}^K F_n^k l_{nj}^{k,f},$$

¹⁷The MRIO table identifies only the composition of production labor across sectors, $\chi_n^s \equiv L_n^s / \sum_{s=1}^K L_n^s$, rather than $v_n^s \equiv L_n^s / L_n$ relative to total city employment (which also includes search and entry labor). We recover v_n^s from χ_n^s using the goods and labor market clearing conditions; the mapping is derived in Appendix C.4.

¹⁸This reinterpretation is consistent with the structural model. The critical mechanism is that lower passenger travel time raises the intensity of procurement-related travel from buyers in n toward potential suppliers in j , thereby improving the distribution of import prices that those buyers achieve. Search provides a tractable microfoundation for how passenger travel time enters sourcing costs and equilibrium trade shares. In this sense, it stands in for a broader class of travel-mediated procurement activities that improve the achieved price distribution, including supplier evaluation, quality verification, negotiation, and relationship maintenance. Nevertheless, this reinterpretation still abstracts from non-procurement business travel.

where $\phi_{nj} \equiv V_{nj}^U / \sum_{k=1}^K F_n^k l_{nj}^{k,f}$ is the ratio of residual to procurement-related travel. Therefore, we have

$$\Delta \ln \frac{V_{nj}}{\sum_{k=1}^K \pi_{jn}^k D_n^k} = -\alpha^P \Delta \ln t_{nj} + \Delta \delta_n + \Delta \epsilon'_{nj}, \quad (36)$$

with $\Delta \epsilon'_{nj} \equiv \Delta \ln(1 + \phi_{nj}) + \Delta \epsilon_{nj}$. A consistent estimate therefore requires that $\Delta \ln(1 + \phi_{nj})$ be orthogonal to $\Delta \ln t_{nj}$ after conditioning on buyer-city fixed effects.

The most plausible violations of this condition attenuate rather than inflate α^P . First, if procurement-related and residual volumes respond to travel time in the same proportion, ϕ_{nj} is invariant to t_{nj} , so that $\Delta \ln(1 + \phi_{nj})$ does not covary with $\Delta \ln t_{nj}$. Equation (36) then estimates α^P consistently. Second, if residual travel is less time-elastic than procurement-related travel, which is plausible given residual travel contains some time-inelastic trips highlighted above, a decline in t_{nj} raises procurement travel by more than residual travel and attenuates the composite elasticity of observed volume toward zero. The estimated α^P is therefore conservative for the structural elasticity that governs sourcing frictions.

Because city-level trade shares are available only in 2012 and 2015, we construct the dependent variable for these two years. We further restrict the sample to intercity pairs, because travel volumes are inferred from UnionPay card transactions and therefore do not measure intra- and intercity travel on a comparable scale.¹⁹ As reported in Table 5, we begin in Columns (1) and (2) by estimating cross-sectional regressions for 2012 and 2015, thereby leveraging the spatial variation in bilateral travel times. In Column (3), we estimate the first-difference specification in Equation (36). Column (3) employs the full sample of city pairs, while Column (4) restricts the estimation sample by excluding pairs with a straight-line distance exceeding 1,200 km. Consistent with the robustness checks conducted for the leisure travel gravity estimation, this spatial restriction mitigates potential biases arising from unobserved air travel substitution. Across all columns, the estimated coefficients remain close and statistically significant. For our subsequent counterfactual simulations, we adopt the point estimate from Column (3), setting our baseline parameter to $\alpha^P = 0.3$.

4.2.4 Trade gravity

Finally, we estimate the trade gravity equation to identify the convexity of search costs, κ . Regarding data availability, the city-level MRIO tables estimate the trade shares for 42 sectors in 2012 and 2015. Therefore, we take the logarithm of the trade share equation (i.e., Equation (16)) and construct the following first-difference specification:

$$\Delta \ln \pi_{jn,t}^k = \frac{\alpha^P}{1 - \kappa} \Delta \ln t_{nj} + \Delta \delta_j^k + \Delta \delta_n^k + \Delta \epsilon_{jn}^k, \quad (37)$$

where Δt_{nj} is the observed change in the minimal travel time from buyer city n to supplier city j across all modes of transport (i.e., HSR, conventional train, and highway), sector-supplier-city

¹⁹For example, overnight lodging and other travel-related spending are much less likely to be recorded when a trip remains within the same city, so intra-city volumes are systematically underestimated. Unlike the leisure-travel gravity, where intra-city trips are excluded from the choice set by construction ($m \neq n$), the search block treats own-city procurement as one destination among others. However, we drop these local pairs to avoid mixing incomparable measurements.

Table 5: The estimation of the business travel volume equation

	Dependent variable: $\ln(V_{nj}/\sum_{k=1}^K \pi_{jn}^k D_n^k)$ or $\Delta \ln(V_{nj}/\sum_{k=1}^K \pi_{jn}^k D_n^k)$			
	2012	2015	First-difference	
	(1)	(2)	(3)	(4)
$\ln t_{nj}$	-0.242*** (0.057)	-0.353*** (0.053)		
$\Delta \ln t_{nj}$			-0.298** (0.129)	-0.370*** (0.123)
Buyer-city fixed effects	Yes	Yes	Yes	Yes
Distance restriction	No	No	No	$\leq 1,200$ km
# of observations	54,976	51,676	51,676	33,262

Notes: This table reports the estimation results for the business travel volume equation. Columns (1) and (2) report cross-sectional estimates of Equation (36) for 2012 and 2015, respectively, where the dependent variable is $\ln(V_{nj}/\sum_{k=1}^K \pi_{jn}^k D_n^k)$. Columns (3) and (4) estimate the first-differenced specification between 2012 and 2015. Column (4) restricts the estimation sample by excluding city pairs with straight-line distances exceeding 1,200 km. Standard errors are two-way clustered by supplier and buyer city levels and reported in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

fixed effects $\Delta \delta_j^k \equiv \kappa/(\kappa - 1) \Delta \ln T_j^k - \kappa \theta^P/(\kappa - 1) \Delta \ln c_j^k$ absorbs the changes in supplier's average productivity and production cost, sector-buyer-city fixed effects $\Delta \delta_n^k$ absorbs the change in the multilateral resistance denominator, and the error term $\Delta \epsilon_{jn}^k$ contains the changes in freight trade costs $\Delta \ln \bar{d}_{jn}$.

The main identification concern is that changes in passenger travel time need not be orthogonal to the freight-cost component in the error term. In particular, improvements in highway infrastructure can simultaneously reduce passenger travel time and the physical cost of shipping goods. To isolate passenger-time variation, we instrument $\Delta \ln t_{nj}$ using the change in travel time purely generated by HSR between 2012 and 2015.²⁰ The exclusion restriction is motivated by the fact that HSR carries passengers rather than freight. HSR-induced travel-time savings directly lower the passenger friction relevant for business search without mechanically lowering the physical cost of shipping intermediate goods. Therefore, it rules out the most immediate correlation between the instrument and $\Delta \ln \bar{d}_{jn}$.

Column (1) of Table 6 reports the IV estimation results. The coefficient is statistically significant and negative, suggesting $\kappa > 1$ which is consistent with our parameter assumption in the model. Furthermore, we exclude city pairs that were already connected by HSR in or before 2012 in Column (2). This restriction removes pairs whose treatment predates the baseline year and makes the identifying comparison more transparent: variation comes from HSR connection improvement between 2013 and 2015 relative to pairs not affected during this period. The coefficient decreases in magnitude but remains statistically significant.

A remaining concern is that the HSR instrument may reflect coordinated corridor-level infras-

²⁰For example, denote the minimal travel time across all modes of transport in 2012 as $t_{nj,2012}$ and the travel time by HSR in 2015 as $t_{nj,2015}^{\text{HSR}}$. The instrument is then defined as $\ln t_{nj,2015}^{\text{HSR}} - \ln t_{nj,2012}$. For city pairs without HSR connection in 2015, the instrument is set to 0.

structure planning. China’s national transport plans organized railway and highway within common multimodal corridors and explicitly treated direct-administered municipalities, provincial capitals, and municipalities with independent planning status as core nodes.²¹ As a result, HSR-induced changes in passenger travel time may remain correlated with changes in residual freight costs through simultaneous investments in highway or other freight infrastructure. To alleviate this concern, we directly control for the change in freight travel time by highway in Column (3). This variable proxies for $\Delta \ln \bar{d}_{jn}$ and therefore identifies the passenger-search channel holding measured freight conditions constant.

As a complementary network-placement sensitivity check, Column (4) excludes city pairs where both the origin and the destination fall into any of the following categories: direct-administered municipalities, provincial capitals, or municipalities with independent planning status, which are particularly exposed to the endogenous network placement. Reassuringly, both estimates in Columns (3) and (4) remain very close to that in Column (2). Although this restriction does not claim quasi-random assignment among the remaining pairs, the stability of the coefficient provides a robustness check, suggesting that our estimates are likely not confounded by this source of endogeneity. For our counterfactual simulations, we adopt the point estimate from Column (4), which suggests that $\alpha^P/(1 - \kappa) = -0.22$. Combining this with the estimate of $\alpha^P = 0.3$, we obtain $\kappa = 2.36$. Finally, with the calibrated $\kappa\theta^P/(\kappa - 1) = 5.94$, we obtain $\theta^P = 3.42$. This implies that, relative to classical trade models without information frictions, the elasticity of trade with respect to freight transport costs is amplified by a factor of 1.74. Additionally, passenger travel costs τ_{nj} enter trade shares with elasticity $1/(\kappa - 1) \approx 0.73$, approximately 12% of the elasticity with respect to freight transport costs. Although this relative magnitude is again modest, the HSR-induced reduction in passenger travel costs is large enough to generate substantial spatial effects.

To conclude this section, we summarize all the calibrated and estimated parameters in Appendix Table E.2.

5 Counterfactual effects of HSR network construction

We now use the estimated model to quantify the aggregate effects of HSR network construction between 2008 and 2015. The EHA approach requires an observed equilibrium as the baseline. Due to data availability, we construct this equilibrium from the spatial distribution of labor in 2010, leisure travel shares in 2011, and sectoral trade shares in 2012. Taken together, we date this composite equilibrium as 2011 and interpret it as characterizing the Chinese spatial economy at the beginning of the 2010s.

The exogenous shock is the reduction in passenger travel time induced by the HSR network. For

²¹For instance, the *Medium- and Long-Term Comprehensive Transport Network Development Plan* (2007) mandates that national multimodal transport corridors, defined as integrated networks comprising multiple transportation modes, connect all direct-administered municipalities, provincial capitals, and municipalities with independent planning status, alongside other major cities. The official document is available at <https://www.ndrc.gov.cn/xxgk/zcxfb/ghwb/200906/W020190905497521610436.pdf> (accessed July 2026).

Table 6: The estimation of the trade gravity equation

	Dependent variable: $\Delta \ln \pi_{jn}^k$			
	(1)	(2)	(3)	(4)
$\Delta \ln t_{nj}$	-0.372*** (0.125)	-0.223* (0.118)	-0.225* (0.118)	-0.220* (0.121)
$\Delta \ln t_{nj}^{\text{freight, highway}}$			-0.807** (0.386)	-0.814** (0.387)
Supplier-city $\times$ sector fixed effects	Yes	Yes	Yes	Yes
Buyer-city $\times$ sector fixed effects	Yes	Yes	Yes	Yes
Excl. pairs connected in or before 2012	No	Yes	Yes	Yes
Excl. core-to-core pairs	No	No	No	Yes
First-stage F-statistic	443,522	428,736	439,564	423,735
# of observations	3,817,884	3,708,474	3,708,474	3,673,908

Notes: This table presents the IV estimation results of Equation (37). Each observation corresponds to a directed buyer–supplier city pair in a given sector. In all columns, $\Delta \ln t_{nj}$ is instrumented by the change in bilateral passenger travel time generated solely by the HSR network construction. Column (1) uses all intercity pairs across the 42 sectors. Column (2) excludes city pairs that were connected by HSR in or before 2012. Column (3) additionally controls for the change in highway freight travel time (see Appendix B.2 for the variable construction). Column (4) further excludes city pairs where both the origin and the destination fall into any of the following categories: direct-administered municipalities, provincial capitals, or municipalities with independent planning status. Standard errors are two-way clustered by supplier and buyer city levels and reported in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

each year t and city pair from n to j , define

$$\hat{t}_{njt} = \frac{t_{njt}^{\text{HSR}}}{t_{nj}^{\text{no-HSR}}}, \quad (38)$$

where t_{njt}^{HSR} is the minimum door-to-door HSR travel time in each year and $t_{nj}^{\text{no-HSR}}$ is the minimum travel time by highway or conventional rail in the baseline year (2011). We set $\hat{t}_{njt} = 1$ for unconnected pairs (including local pairs) and whenever the ratio exceeds one, so that HSR never lengthens travel time.²²

Because some lines had already opened in the early 2010s, the observed allocations already embed part of the treatment. To recover a no-HSR baseline, we apply the inverse shock $\{1/\hat{t}_{nj,2011}\}_{n,j \in \mathcal{N}}$ and solve the equilibrium in relative changes. The resulting allocation preserves the fundamentals of the observed early-2010s economy while shutting down the HSR network. All subsequent counterfactuals introduce actual HSR travel-time reductions relative to this no-HSR baseline. On the production side, the counterfactuals are implemented with three sectors—primary, secondary, and tertiary—obtained by aggregating the 42 MRIO industries, as detailed in Appendix B.5. Finally, we compute the counterfactual equilibrium for each year from 2008 to 2015 relative to the no-HSR baseline. Because the model is static, each of these equilibria is a long-run spatial allocation given that year’s HSR network, holding baseline fundamentals fixed. It is not the short-run impact of the

²²For a small number of door-to-door travels, highway remains faster than HSR. In these cases, we treat HSR as having no effect on travel time.

network in that calendar year, nor a dynamic transition between years.

5.1 Passenger travel-time reductions

We start by examining how HSR affects average travel times through both direct time reductions and the endogenous reallocation of travel across destinations in the counterfactual equilibrium. Specifically, define the national average travel time as

$$\bar{t} \equiv \frac{\sum_{n=1}^N \sum_{j=1}^N t_{nj} V_{nj}}{\sum_{n=1}^N \sum_{j=1}^N V_{nj}} = \sum_{n=1}^N \sum_{j=1}^N t_{nj} s_{nj}, \quad (39)$$

where V_{nj} denotes the travel volume (either business or leisure) from city n to city j and $s_{nj} \equiv V_{nj} / \sum_{n=1}^N \sum_{j=1}^N V_{nj}$ is the share of total volume for pair from n to j . Let $S_{nj} \equiv t_{nj} s_{nj} / \sum_{n=1}^N \sum_{j=1}^N t_{nj} s_{nj}$ denote the corresponding share of total travel time. The change in national average travel time is then

$$\hat{t} = \sum_{n=1}^N \sum_{j=1}^N S_{nj} \hat{t}_{nj} \hat{s}_{nj},$$

which decomposes as

$$\hat{t} - 1 = \underbrace{\sum_{n=1}^N \sum_{j=1}^N S_{nj} (\hat{t}_{nj} - 1)}_{\text{direct time effects}} + \underbrace{\sum_{n=1}^N \sum_{j=1}^N S_{nj} (\hat{s}_{nj} - 1)}_{\text{composition effects}} + \underbrace{\sum_{n=1}^N \sum_{j=1}^N S_{nj} (\hat{s}_{nj} - 1) (\hat{t}_{nj} - 1)}_{\text{interaction effects}}. \quad (40)$$

The first term is the mechanical contribution of the HSR-induced time reductions, holding travel shares fixed at their baseline values. The second term captures shifts in the composition of travel toward routes that were already longer or shorter in the baseline. The third term reflects the covariance between time reductions and share changes. Because travelers reallocate toward destinations whose times have fallen, this interaction is expected to be negative. We measure baseline shares $\{s_{nj}, S_{nj}\}_{n,j \in N}$ from 2011 UnionPay volumes and evaluate all changes relative to the no-HSR baseline. Appendix D.1 details the construction of $\hat{s}_{nj}$ for business and leisure travel.

The two types of travel enter the decomposition differently. For business travel, the model delivers procurement-related volumes from the search block (Equation (34)), so all three terms in Equation (40) are well defined at the national level. We hold the residual-to-procurement ratio ϕ_{nj} fixed, so that differential responses of residual travel are shut down. As a result, our calculation quantifies the change in average travel time when residual business travel scales with procurement. For leisure travel, the model determines destination shares conditional on outbound traveling, but not the number of intercity travelers originating in each city. National composition and interaction effects are therefore not identified for leisure travel. Direct time effects remain well defined for both, because they depend only on baseline time shares and the HSR time shocks.

Panel A of Table 7 reports the 2015 cumulative effects. National average business travel time falls by 18.406%. Direct time effects account for a dominant portion of this change. The two endogenous terms are opposite in sign: composition effects of 2.429% and interaction effects of

−4.071%. The positive composition term means that procurement travel shifts toward routes that were already more time-intensive. The negative interaction term, as expected, indicates that these same routes are the ones whose times decline. The two terms largely offset. Leisure travel times fall in parallel and the national direct time effect is −15.664%. We graphically show the year-by-year decomposition of average business travel time in Appendix Figure E.4.

These time reductions are unevenly distributed. Applying the city-level analogue of Equation (39), we compute the change in average travel time for each city and summarize the cross-city distribution in Panel B of Table 7. Appendix Figures E.5 and E.6 map the 2015 city-level changes in average business and leisure travel time. Furthermore, we classify a city as large if its population in the no-HSR baseline lies above the median across all cities. Restricting to cities connected by HSR in 2015, Panel A shows that large cities experience systematically larger direct time effects than small cities.²³ Therefore, even among connected cities, the HSR shock delivers unequal direct treatment exposure. Our subsequent analysis shows that this exposure heterogeneity is the primary driver of the spatial redistribution.

Table 7: Percentage change in average travel time in 2015 by business and leisure travel

	Business travel	Leisure travel
<i>Panel A. National average</i>		
Total effects	−18.406	
Direct time effects	−16.763	−15.664
Large cities with HSR	−19.580	−18.423
Small cities with HSR	−11.228	−11.059
Composition effects	2.429	
Interaction effects	−4.071	
<i>Panel B. Effects by city</i>		
Mean	−7.289	−10.701
Maximum	3.101	2.627
Minimum	−41.072	−56.909

Notes: This table reports the 2015 counterfactual percentage change in average travel time relative to the no-HSR baseline. Panel A decomposes the national average using Equation (40). Business travel is measured by procurement-related search volumes, with residual business travel scaled in proportion. Leisure-travel composition, interaction, and total effects are not identified nationally, because the model does not determine the number of intercity leisure travelers by origin. A city is classified as large if its population in the no-HSR baseline lies above the median across all cities, and as having HSR if at least one of its directed pairs receives a travel-time reduction in 2015. Direct time effects for large and small cities with HSR restrict that calculation to trips originating in each group. Panel B reports the unweighted distribution of city-level percentage changes in average travel time. Baseline shares are measured from 2011 UnionPay volumes.

5.2 Spatial reallocation of labor

We next investigate how HSR-induced travel-time reductions reshape the long-run spatial allocation of labor. In our counterfactual simulations, workers adjust their residential locations in response

²³Recall that we set $\hat{t}_{njt} = 1$ whenever HSR does not reduce door-to-door travel time relative to highway or conventional rail. Accordingly, throughout the counterfactual analysis, we classify city n as connected by HSR in year t if $\hat{t}_{njt} < 1$ for at least one directed pair involving an outbound trip from city n . This model-based classification may differ slightly from whether the city has HSR service in the actual network.

to both the direct improvements in leisure travel accessibility facilitated by the HSR network and the ensuing general equilibrium effects on real wages. Figure 1 plots the changes in equilibrium migration stocks induced by the HSR network across different years. Furthermore, because HSR alters the spatial equilibrium through two distinct mechanisms (namely, the leisure travel and business search channels), we decompose the aggregate migration response into these two margins by muting the relative changes in travel times for each respective channel at a time.²⁴

Relative to the no-HSR baseline, the 2015 network increases the stock of intercity migrants by 6.637%. The business-search channel alone raises it by 1.742%, whereas the leisure channel alone raises it by 4.351%. Their interaction contributes a further 0.544 percentage points. Thus, the leisure channel generates a substantially larger migration response than the production-side search channel. This response arises because the HSR network improves leisure accessibility unevenly across cities. Workers' residential choices depend on relative accessibility: an improvement shared proportionately by all cities raises the utility from leisure travel broadly, but does little to alter their relative attractiveness. By contrast, the historical rollout disproportionately improves accessibility in particular cities and thereby induces workers to reallocate across space.

This distinction also explains the non-monotonic migration response near the end of the sample period. Specifically, the leisure-only effect reaches 4.861% in 2014 but falls to 4.351% in 2015. The decline is consistent with the further network expansion in 2015 spreading accessibility improvements more evenly across cities, thereby mitigating the spatial asymmetries that previously generated strong migration incentives. Therefore, a larger HSR network does not necessarily induce greater aggregate migration; rather, the magnitude of labor reallocation hinges on how marginal network expansions alter relative accessibility across regions.

The unequal travel-time reductions documented in the previous subsection provide a natural lens through which to interpret the spatial direction of this reallocation. Large cities with HSR receive substantially larger reductions in both business and leisure travel times than small cities with HSR. These uneven shocks translate into unequal accessibility improvements and, consequently, unequal changes in cities' relative attractiveness as residential locations. Consistent with this pattern, Panel A of Table 8 shows that the resident population of large connected cities increases by 0.967%, equivalent to approximately 6.807 million workers. In contrast, the population of small connected cities is nearly unchanged, increasing by only 0.099%, or 0.131 million workers. Cities without HSR lose 1.602% of their population, or approximately 6.937 million workers. Hence, almost all of the net population drawn into connected cities accrues to large cities. We show the change in the resident population in 2015 by city in Appendix Figure E.7.

Panel B further presents the origin of this spatial concentration. The number of workers with *hukou* in small connected cities who reside in large connected cities increases by approximately 0.489 million. The reverse allocation from large-city *hukou* origins to small connected cities also increases, but by only 0.068 million. This directional asymmetry resembles the siphoning pattern from peripheral locations toward larger urban centers documented in the HSR literature (e.g., Koster

²⁴In our counterfactual non-linear system expressed in exact hat algebra (Definition A.1 in Appendix A.5), muting the leisure travel channel corresponds to setting $\hat{t}_{nj} = 1$ for all n and j in Equation (A.4), whereas muting the business search channel is equivalent to setting $\hat{t}_{nj} = 1$ in Equation (A.5).




Figure 1: Change in the number of total migrants due to HSR over time

Notes: This figure reports the percentage change in the equilibrium stock of intercity migrants under each annual HSR network relative to the no-HSR baseline. “Both channels” allows HSR-induced travel-time reductions to affect both leisure travel and business search. The search-only and leisure-only counterfactuals hold the other friction at its baseline level. The synergy effect is the full effect minus the two channel effects.

et al., 2022). Nevertheless, the small-to-large reallocation is not the dominant source of population growth in large cities. Workers originating from cities without HSR account for an increase of 5.654 million, approximately 83% of the total population gain in large connected cities. The near-zero population change of small connected cities similarly masks substantial offsetting changes. A decline of 0.301 million in residents originating from small connected cities is offset by increases of 0.364 million from cities without HSR and 0.068 million from large connected cities. Thus, the stable aggregate population of small connected cities conceals a considerable change in its composition.

Taken together, the actual HSR rollout induces a clear concentration of labor in large connected cities. However, because these cities also receive substantially larger travel-time reductions, the group-level contrast does not identify an independent effect of city size. It combines differences in treatment exposure with any differences in responses conditional on exposure. We further separate these two margins in the section about heterogeneity.

5.3 Aggregate effects on real wages and welfare

We now turn from the spatial reallocation of labor to the associated changes in real income and welfare. The national average real wage aggregates local consumption wages and therefore records the production-side consequences of HSR, including the reallocation of labor across cities. Since labor is the only factor that generates value added in goods production, the national average real wage can be interpreted as the change in real GDP per worker as well. Expected utility further

Table 8: Spatial reallocation of labor in 2015

	Residence city group		
	Large cities with HSR	Small cities with HSR	Cities without HSR
<i>Panel A. Change in total resident population</i>			
Percentage change	0.967	0.099	−1.602
Level change (million persons)	6.807	0.131	−6.937
<i>Panel B. Level change by hukou-origin group</i>			
Large cities with HSR	0.664	0.068	−0.731
Small cities with HSR	0.489	−0.301	−0.189
Cities without HSR	5.654	0.364	−6.017

Notes: Panel A reports changes in total population by residence-city group. Panel B disaggregates the resident population by *hukou*-origin and reports its level change in million persons. These quantities compare group-level equilibrium stocks and should not be interpreted as longitudinal transition rates for individuals residing in a particular group in the baseline. A city is classified as large if its population in the no-HSR baseline lies above the median across all cities, and as having HSR if at least one of its directed pairs receives a travel-time reduction in 2015.

incorporates leisure accessibility and residential choice.

Define the national average real wage as

$$\bar{w} \equiv \frac{\sum_{n=1}^N w_n L_n / P_n}{\sum_{n=1}^N L_n}. \quad (41)$$

With aggregate population fixed, the relative change is

$$\hat{\bar{w}} = \sum_{n=1}^N \frac{\hat{w}_n}{\hat{P}_n} \hat{L}_n \varsigma_n,$$

where $\varsigma_n \equiv (w_n L_n / P_n) / (\sum_{n=1}^N w_n L_n / P_n)$ is a weight that measures city n 's baseline share of the national real wage bill.²⁵

Figure 2 reports the implied changes under each annual HSR network. Relative to the no-HSR baseline, the 2015 network raises the national average real wage by 2.185%. The business-search and leisure channels contribute 1.092% and 0.951%, respectively, with a small synergy of 0.142 percentage points. The search-only effect rises with network expansion, consistent with a production-side shock whose first-order effect is to lower sourcing costs as more pairs become cheaper to search.

The leisure travel channel is not a production-side shock, yet it raises the national average real wage by nearly as much as the search channel. It does so because the historical network improves accessibility unevenly and, as documented above, draws workers into large connected cities. Those destinations have higher baseline real wages, so the reallocation raises $\hat{\bar{w}}$ in part through composition. The leisure-only effect peaks at 0.972% in 2014 and falls slightly to 0.951% in 2015, paralleling the migration response previously found: once further expansion spreads accessibility gains more

²⁵Baseline L_n is the model's labor allocation, constructed from census migration shares, and w_n is recovered from goods and labor market clearing. Baseline P_n does not enter clearing conditions, so it is not recovered. We use the observed corresponding provincial price index in 2011 for the construction of ς_n .

evenly, the incentive to reallocate labor weakens. In this sense, a consumption amenity that is itself orthogonal to production can still raise real GDP when the underlying travel-time shock is spatially concentrated.




Figure 2: Change in the national average real wage due to HSR over time

Notes: This figure reports the percentage change in the national average real wage under each annual HSR network relative to the no-HSR baseline. “Both channels” allows HSR-induced travel-time reductions to affect both leisure travel and business search. The search-only and leisure-only counterfactuals hold the other friction at its baseline level. The synergy effect is the full effect minus the two channel effects.

We measure welfare as the expected utility of workers grouped by *hukou* origin. For each origin j , taking the maximum of $U_{jn}(i)$ in Equation (4) over residential destinations n and then integrating over the Fréchet residential shocks yields the origin-specific expected utility proportional to Θ_j^{1/θ^R} , where Θ_j is the denominator of the migration share in Equation (5). Using the local share to express Θ_j , the relative change for origin j is

$$\hat{U}_j = \frac{\hat{w}_j}{\hat{p}_j} \left(\hat{\Phi}_j^T \right)^{\psi/\theta^T} \hat{\lambda}_{jj}^{-1/\theta^R}. \quad (42)$$

We summarize the national change by the *hukou*-population-weighted geometric mean:²⁶

$$\hat{U} = \prod_{j=1}^N \left(\hat{U}_j \right)^{\bar{L}_j/\bar{L}}. \quad (43)$$

Taking logs and substituting Equation (42) decomposes the change into local real wages, leisure

²⁶The geometric mean is used so that $\ln \hat{U}$ equals the *hukou*-weighted sum of $\ln \hat{U}_j$ and therefore decomposes exactly. When origin-specific changes are small, it is first-order equivalent to the corresponding arithmetic mean.

accessibility, and the value of migration:

$$\ln \hat{U} = \sum_{j=1}^N \frac{\bar{L}_j}{\bar{L}} \ln \frac{\hat{w}_j}{\hat{P}_j} + \sum_{j=1}^N \frac{\bar{L}_j}{\bar{L}} \frac{\psi}{\theta^T} \ln \hat{\Phi}_j^T + \sum_{j=1}^N \frac{\bar{L}_j}{\bar{L}} \left(-\frac{1}{\theta^R} \ln \hat{\lambda}_{jj} \right). \quad (44)$$

Figure 3 presents the total welfare effects and the decomposition. In total, the HSR network in 2015 leads to a 3.14% increase in national welfare, where the accessibility component accounts for 77% of this increase. The accessibility term accounts for most of the welfare gain because HSR generates large, uneven changes in leisure travel accessibility. Although the migration elasticity with respect to accessibility is only about one-tenth of that with respect to real wages, these accessibility changes are large enough to dominate $\ln \hat{U}$.

The real wage component is not the national average real wage in Equation (41). In 2015 it is only 0.338%, compared with a 2.185% increase in $\hat{w}$. Two differences account for the gap. First, the real wage component contains no reallocation term $\hat{L}_n$, so it does not record workers moving toward cities with higher real wages. Second, it averages each origin's own-city real wage change by *hukou* population. Those weights are fixed at origin endowments and therefore put more mass on cities experiencing labor outflows than do residence weights. Conversely, the welfare improvement through labor reallocation is captured by the migration component.

Muting each travel-time shock in turn yields the same ranking. Under the 2015 HSR network, the leisure channel alone raises national welfare by 2.574%, and the search channel alone by 0.554%, with negligible synergy. The two channels are therefore far more unbalanced for welfare than for the national average real wage. The year-by-year channel decomposition is reported in Appendix Figure E.8.

Additionally, we graphically show each city's changes in the real wage and *hukou*-residence welfare under the 2015 HSR network in Appendix Figure E.9 and will further discuss this heterogeneity in the subsequent analysis.

To isolate the contribution of endogenous supplier search, Appendix D.2 compares the full model with a no-search benchmark that replaces the search block with standard Eaton–Kortum sourcing while retaining the same baseline allocation and all other spatial-equilibrium mechanisms. Because HSR changes passenger travel costs but not freight costs, this benchmark provides no direct channel through which HSR affects firms' sourcing decisions. Its remaining real-wage gain of 0.71% is generated by leisure-induced migration and its propagation through standard Eaton–Kortum trade adjustment. In the full model, the same leisure shock raises the national average real wage by 0.951%, because the induced changes in wages and production costs additionally feed back into firms' search intensity. Activating the direct business-search channel further raises the total gain to 2.185%, suggesting that endogenous supplier search is central to the real-income effects of HSR. The appendix also reports fixed-migration and fixed-trade diagnostics that trace how labor reallocation and trade-share adjustment propagate these effects.

Furthermore, as a robustness check, we set $\hat{t}_{nj} = 1$ for pairs more than 1,200 km apart, so that HSR is not permitted to reduce travel time on distances at which air travel is a relevant substitute. As shown in Appendix Figure E.10, the national real-wage and welfare paths remain close to the

baseline series, with a modest attenuation in later years.

Finally, we examine the sensitivity of the aggregate effects to the calibrated trade elasticity with respect to freight costs. As discussed in Section 4, the benchmark value of 5.94 for $\kappa\theta^P/(\kappa-1)$ may overstate the true composite elasticity. We therefore re-solve the 2015 counterfactual with smaller values of 5.44 , 4.94 , 4.44 , and 3.94 . As reported in Appendix Table E.3, the national real-wage gain remains stably between 2.17% and 2.33% . By contrast, the welfare gain increases moderately from 3.14% to 3.61% as the elasticity declines. Our benchmark calibration is therefore, if anything, conservative along this dimension.




Figure 3: Change in the national welfare due to HSR over time

Notes: This figure reports the log change in national welfare under each annual HSR network relative to the no-HSR baseline, together with the three components in Equation (44). National welfare is the *hukou*-population-weighted geometric mean of origin-specific expected utilities. The real-wage, accessibility, and migration components sum to the total.

6 Spatial heterogeneity

The aggregate gains from HSR conceal substantial heterogeneity across cities. The preceding analysis shows that the HSR network reallocates population toward large connected cities, while cities outside the network lose population. These group-level differences, however, combine two conceptually distinct margins. Cities may receive different reductions in passenger travel time, and they may respond differently to a given reduction because of differences in their baseline economic structure.

This section separates these margins. We first examine whether the larger gains of large connected cities reflect greater treatment exposure or an intrinsically stronger response to HSR. We then study

how baseline trade and migration linkages shape the conversion of direct HSR exposure into local economic activity. Finally, we examine cities without direct HSR exposure and quantify how they are affected indirectly through their migration destinations and export markets.

6.1 City size, treatment exposure, and local gains

We begin by summarizing the counterfactual effects of the 2015 HSR network across three groups of cities in Table 9. Aggregating across residence cities within each group, local real GDP, defined as the group sum of $w_n L_n / P_n$, increases by 4.486% in large connected cities, compared with 1.648% in small connected cities, while it falls by 2.73% in cities without direct HSR exposure. Both components of local real GDP exhibit the same ordering. Large connected cities experience increases of 0.967% in resident population and 3.485% in average real wages, whereas the corresponding increases in small connected cities are 0.099% and 1.548%. Cities without HSR experience declines along both margins.

Welfare gains are less uneven across connected-city groups. Workers whose *hukou* origins are in large and small connected cities experience welfare gains of 5.391% and 4.439%, respectively. This partly reflects the direct positive contribution of improved leisure accessibility to welfare. The search channel, by contrast, affects local real wages through adjustments in sourcing and trade, which reallocate demand across cities and make real-wage gains more dependent on a city's exposure relative to the rest of the network. It also reflects the origin-based welfare measure: gains accruing to migrants are assigned to their *hukou*-origin groups rather than to their counterfactual residence cities. Workers originating in cities without HSR nevertheless experience a small average welfare loss of 0.475%.²⁷

Table 9: Spatial heterogeneity in the effects of the 2015 HSR network

	Large cities with HSR	Small cities with HSR	Cities without HSR
Local real GDP	4.486	1.648	-2.730
Resident population	0.967	0.099	-1.602
Real wage	3.485	1.548	-1.147
Welfare (<i>hukou</i> origin)	5.391	4.439	-0.475

Notes: This table reports changes under the 2015 HSR network relative to the no-HSR baseline, with both the business-search and leisure-accessibility channels active. The first three rows group cities by residence. Local real GDP is the group sum of $w_n L_n / P_n$, and real wage is group local real GDP divided by resident population. Welfare is the *hukou*-population-weighted mean log change in origin-specific expected utility, so welfare groups refer to *hukou* origins rather than residences. A city is large if its no-HSR-baseline population exceeds the sample median, and is classified as having HSR if at least one outbound directed pair experiences a travel-time reduction in 2015.

These group-level patterns relate to a longstanding question about whether HSR reinforces large urban centers at the expense of smaller connected locations. Existing studies reach different conclusions about the effects of HSR on smaller cities (e.g., Lin, 2017; Ahlfeldt and Feddersen,

²⁷Appendix Table E.4 repeats this comparison when HSR operates only through business search or only through leisure accessibility. The ranking across the three city groups remains the same under both counterfactuals. Thus, the concentration of gains in large connected cities is not driven exclusively by either channel, although the relative contributions of the two channels differ across outcomes.

2018; Koster et al., 2021). A limitation of a binary comparison between connected and unconnected cities is that connection does not imply a common treatment. The economic exposure of a city depends on which destinations become more accessible, how much travel time falls, and the baseline importance of the affected routes.

To separate treatment intensity from conditional responses, we define city n 's direct HSR exposure, E_n , as the direct reduction in its average business travel time holding baseline travel shares fixed. Equivalently, E_n is the negative of the city-level counterpart to the direct time effect in Equation (40). Unlike the total change in average travel time, this measure excludes endogenous destination reallocation and captures the extent to which the realized HSR network mechanically favored each city. As reported in Table 7, the direct reduction in average business (leisure) travel time is 19.58% (18.423%) for large connected cities and 11.228% (11.059%) for small connected cities. Large connected cities therefore receive substantially greater treatment exposure through both travel channels.

We then examine whether large and small cities respond differently conditional on E_n . Figure 4 plots each city's change in local real GDP against its direct HSR exposure separately by city-size group. Local real GDP increases strongly with the reduction in travel time in both groups. More importantly, the two fitted gradients are similar. Small connected cities do not systematically contract when they receive substantial HSR exposure, nor do large cities experience larger gains when exposure is held constant.²⁸

Therefore, large connected cities gain more local economic activity primarily because the historical network delivers larger travel-time reductions to them, rather than because large cities possess an intrinsically stronger response to a given reduction. Appendix Figure E.11 shows that the corresponding gradients for resident population and real wages yield the same qualitative conclusion.

6.2 Baseline spatial linkages and local gains

City size is not the only characteristic that may shape the incidence of HSR. Even among cities receiving similar direct exposure, baseline spatial linkages can determine how the shock propagates through production networks and labor markets. We therefore examine which linkages amplify or dampen the conversion of E_n into local real GDP by running the following regression across connected cities:

$$\widehat{RealGDP}_n - 1 = \alpha + \beta E_n + \delta S_n + \gamma (E_n \times S_n) + \epsilon_n, \quad (45)$$

where S_n measures the baseline spatial linkages.

We primarily focus on the following four variables. On the trade side, the local sourcing share is defined as the weighted share of intermediate inputs that city n purchases locally:

$$LocalSourcing_n = \sum_{k=1}^K \pi_{nn}^k \frac{X_n^k}{\sum_{s=1}^K X_n^s},$$

²⁸Although not reported, using the direct reduction in leisure travel time as the exposure measure yields very similar fitted gradients because the city-level business and leisure travel-time reductions are highly correlated.




Figure 4: Direct HSR exposure and local real GDP by city size

Notes: This figure plots the model-implied relative change in local real GDP under the 2015 HSR network against the positive direct reduction in city-level average business travel time. The exposure measure is calculated as the city-level counterpart of the direct time effect in Equation (40), holding baseline travel shares fixed. The sample is restricted to cities with positive direct business-travel exposure. A city is large if its population in the no-HSR baseline exceeds the sample median. Dashed lines report group-specific linear fits.

where the weight $X_n^k / \sum_{s=1}^K X_n^s$ reflects the relative importance of sector k in city n . The external-market share measures the importance of city n as a supplier to buyers in other cities:

$$ExternalMarket_n = \sum_{k=1}^K \frac{X_n^k}{\sum_{s=1}^K X_n^s} \sum_{m \neq n} \pi_{nm}^k \frac{X_m^k}{\sum_{m' \neq n} X_{m'}^k}. \quad (46)$$

For each sector, the inner summation measures the share of total purchases by all other cities that is sourced from city n . We then aggregate these sector-specific market shares by the city's sectoral weight. A larger value therefore indicates that city n is a more important supplier in external markets.

On the migration side, λ_{nn} is the share of workers with *hukou* in city n who remain in that city. We refer to it as the local retention rate. We also construct a baseline measure of the city's attractiveness to workers originating elsewhere:

$$InMigration_n = \sum_{j \neq n} \lambda_{jn} \frac{\bar{L}_j}{\sum_{j' \neq n} \bar{L}_{j'}}.$$

All four variables are standardized to have mean zero and unit standard deviation.

Table 10 reports the estimation results. Column (1) first relates local real GDP to direct HSR exposure alone. The estimated coefficient is positive and highly significant, and the specification accounts for approximately 36% of the cross-city variation in the model-implied local real GDP effects. Column (2) adds city size and its interaction with exposure. The interaction coefficient

is close to zero and statistically insignificant, confirming the finding in the previous subsection that large and small connected cities exhibit similar responses conditional on receiving the same travel-time reduction.

Column (3) shows that the interaction between direct exposure and the local sourcing share is positive and statistically significant, whereas the interaction with the external-market share is negative and marginally significant. Taken together, these estimates suggest that cities with more locally oriented baseline production networks experience larger local real GDP gains from a given travel-time reduction. These cities source a larger share of inputs locally and account for a smaller share of demand in external markets before the HSR expansion. The pattern is consistent with larger marginal gains from passenger connectivity where intercity production linkages were initially weaker.

The migration variables reflect a city's baseline capacity to retain and attract workers. The interactions of exposure with both the local retention rate and in-migration attractiveness are positive. This pattern suggests that improved connectivity may generate larger local GDP responses in cities that are already better able to keep local workers and draw workers from elsewhere. The estimates are not sufficiently precise to establish this amplification conclusively, so we view the evidence as suggestive.

Table 10: Baseline spatial linkages and local real GDP responses among connected cities

	Dependent variable: relative change in local real GDP		
	(1)	(2)	(3)
Time reduction	0.358*** (0.039)	0.367*** (0.095)	0.369*** (0.090)
Time reduction × Large city		−0.004 (0.105)	−0.025 (0.100)
Time reduction × Local sourcing share			0.066** (0.031)
Time reduction × External-market share			−0.123* (0.066)
Time reduction × Local retention rate			0.073 (0.044)
Time reduction × In-migration attractiveness			0.102 (0.093)
Main effects of moderating variables	No	Yes	Yes
Adjusted R^2	0.361	0.355	0.494
# of observations	160	160	160

Notes: This table reports cross-city regressions of the model-implied relative change in local real GDP under the 2015 HSR network relative to the no-HSR baseline. The sample contains cities with positive direct business-travel exposure. “Time reduction” is the positive city-level counterpart of the direct time effect in Equation (40), calculated while holding baseline business-travel shares fixed. A city is large if its population in the no-HSR baseline exceeds the sample median. The local sourcing share, external-market share, local retention rate, and in-migration attractiveness are measured in the no-HSR baseline and standardized within the estimation sample. All specifications include a constant. Whenever an interaction between treatment exposure and a moderating variable is included, the moderating variable is also included, though not reported. Heteroskedasticity-robust standard errors are reported in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

6.3 Indirect effects on cities outside the HSR network

Cities without HSR receive no direct reduction in their own average travel time, but they are nevertheless affected in the spatial equilibrium. Table 9 shows that these cities lose both population and local real GDP under the 2015 network. We now examine which unconnected cities experience larger spillover losses.

We construct two measures of indirect network exposure. First, migration-destination exposure is defined as the baseline-migration-weighted HSR exposure of the destinations chosen by workers originating in city n :

$$E_n^{\text{indirect,migration}} = \sum_{m \neq n} \lambda_{nm} E_m.$$

This measure is large when residents of city n have strong baseline migration links to destinations receiving substantial HSR exposure. Second, importer exposure is defined as the baseline-trade-weighted HSR exposure of cities purchasing intermediate inputs from city n :

$$E_n^{\text{indirect,importer}} = \sum_{k=1}^K \frac{X_n^k}{\sum_{s=1}^K X_n^s} \sum_{m \neq n} E_m \pi_{nm}^k \frac{X_m^k}{\sum_{m' \neq n} X_{m'}^k},$$

which further incorporates importer's treatment exposure in Equation (46). A larger $E_n^{\text{indirect,importer}}$ means that city n is more strongly linked to buyers whose HSR exposure increases substantially.

For the sample of cities without direct HSR exposure, we estimate

$$\widehat{RealGDP}_n - 1 = \alpha + \delta Large_n + \beta_1 E_n^{\text{indirect,migration}} + \beta_2 E_n^{\text{indirect,importer}} + \epsilon_n. \quad (47)$$

Table 11 reports the estimation results. Column (1) shows that, in the absence of further controls, large unconnected cities experience a larger decline in local real GDP than small unconnected cities. This size difference largely disappears once the two measures of partner exposure are introduced in Column (2). Both migration-destination exposure and importer exposure have a negative coefficient with statistical significance. Together, these variables explain about 40% of the cross-city variation.

The migration result reflects a spatial pull mechanism. When the destinations to which workers from city n are initially likely to migrate receive larger accessibility and real-income gains, workers have stronger incentives to leave city n . The importer result reflects a demand-diversion mechanism. When the buyers importing from city n become more exposed to HSR, they can search more broadly across alternative supplier locations, reducing demand for inputs produced in the unconnected city. Thus, the cities experiencing larger losses are those whose workers and external demand are more strongly linked to locations that become more accessible.

7 Alternative HSR network expansion

The preceding analysis quantifies the effects of China's historical HSR rollout, whose routes were concentrated among large cities and high-volume corridors. We now examine how the aggregate and spatial effects of HSR depend on the coverage of the network. We first construct a universal network

Table 11: Indirect HSR exposure and local real GDP in cities without HSR

	Dependent variable: relative change in local real GDP	
	(1)	(2)
Large city	−0.010*** (0.003)	0.003 (0.004)
Migration-destination exposure		−0.011*** (0.001)
Importer exposure		−0.004* (0.002)
Adjusted R^2	0.062	0.399
# of observations	142	142

Notes: The dependent variable is the model-implied change in local real GDP under the 2015 HSR network relative to the no-HSR baseline. The sample contains cities receiving no direct travel-time reduction. A city is large if its population in the no-HSR baseline exceeds the sample median. Migration-destination exposure and importer exposure are standardized. Heteroskedasticity-robust standard errors are reported in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

connecting all city pairs and then consider a sequence of alternative networks that progressively add links in descending order of baseline travel volume. These counterfactuals allow us to distinguish gains concentrated in important links from gains that require broader passenger connectivity, and to examine how the expansion of coverage changes the balance between production, welfare, and spatial reallocation.

To assign HSR travel times to city pairs that were not connected in 2015, we use the long-haul relationship estimated in Section 3. Specifically, the predicted HSR travel time between cities n and j is

$$\hat{t}_{nj}^{\text{HSR}} = \exp(0.764 + 0.781 \times \ln \text{Distance}_{nj}).$$

For city pairs already connected in 2015, we retain their observed HSR travel times. For all other pairs, we use the predicted value above. As in the historical counterfactual, we impose no travel-time reduction whenever HSR would be slower than the corresponding no-HSR mode.

7.1 Universal network

We first consider a universal network in which every pair of cities is connected. Table 12 compares its effects with those of the historical network in place by 2015.

Panel A shows that universal connectivity raises the national average real wage by 2.657%, compared with 2.185% under the historical network. The historical network therefore already realizes approximately 82.2% of the real-wage gain under universal connectivity. Connecting all remaining city pairs generates an additional gain of only 0.472 percentage points, despite the substantial increase in network coverage.

The similarity in the total effects conceals a marked change in their composition. In the single-channel counterfactuals, the real-wage effect of business search more than doubles from 1.092% to 2.332%, whereas the effect of leisure accessibility falls from 0.951% to 0.27%. As discussed in Section 5.3, business connectivity directly lowers firms' sourcing costs, while leisure accessibility

affects national real wages indirectly by reallocating workers toward large cities. When accessibility improvements become more evenly distributed, they generate smaller changes in relative residential attractiveness and significantly weaken this reallocation margin.

Panel B examines whether the aggregate gains are accompanied by a greater spatial concentration of economic activity. Let $\text{Gini}(\cdot)$ denote the Gini coefficient across cities. The change in the cross-city Gini of local real GDP is given by

$$\widehat{\text{Gini}}^{\text{GDP}} - 1 = \frac{\text{Gini} \left(\{w_n \hat{w}_n L_n \hat{L}_n / (P_n \hat{P}_n)\}_{n=1}^N \right)}{\text{Gini} \left(\{w_n L_n / P_n\}_{n=1}^N \right)} - 1. \quad (48)$$

This measure treats each city as a geographic unit and captures the concentration of aggregate real economic activity across locations.²⁹

The historical network raises this spatial Gini by 2.387%, whereas the universal network leaves it nearly unchanged, with an increase of only 0.03%. The channel-specific results reinforce this contrast. Under the historical network, both business search and leisure accessibility increase spatial concentration. Under universal connectivity, the search channel instead reduces the Gini by 0.27%, while the leisure channel raises it by only 0.235%.

The welfare comparison is somewhat different. Panel C shows that universal connectivity raises national welfare by 8.146%, compared with 3.140% under the historical network. The historical network therefore realizes only approximately 38.5% of the universal-network welfare gain. The search-only welfare effect rises from 0.554% to 2.155%, while the leisure-only effect rises from 2.574% to 5.984%. Unlike the real-wage effect of leisure accessibility, residents benefit from reaching a broader set of destinations even when comparable improvements occur elsewhere.

Taken together, the historical network captures most of the potential real-wage gain but a relatively smaller share of the potential welfare gain, while generating considerably greater spatial concentration than the universal network. This comparison does not imply that a lower concentration of total output is itself welfare-improving. Rather, it shows that broader coverage can serve a distinct place-based objective by distributing economic activity more evenly across cities. The next subsection examines this trade-off more systematically by varying network coverage according to baseline business-travel volume.³⁰

7.2 Counterfactual rollout under strict travel prioritization

Section 3 shows that China's historical HSR rollout was volume-oriented: high-volume city pairs were disproportionately connected during the early stages of construction. The historical network was not, however, determined solely by passenger demand. Nearly one-third of the top 1% of

²⁹Again, we approximate P_n using a cumulative provincial consumer price index through 2011.

³⁰Appendix D.3 reports a supplementary calculation of the cross-city dispersion of residence-based welfare. Because baseline city amenities are not recovered, the calculation assumes common local amenities and constructs the city-level index using real wages and recovered leisure accessibility. Under this approximation, the historical and universal networks raise the cross-city welfare Gini by 1.363% and 1.236%, respectively. Universal connectivity therefore modestly attenuates spatial welfare dispersion, although we highlight that these results should be interpreted as suggestive.

Table 12: Comparison of the effects of universal and historical HSR networks

	Network in 2015	Universal network
<i>Panel A. Percentage change in national real wage/GDP</i>		
Both channels	2.185	2.657
Search channel	1.092	2.332
Leisure channel	0.951	0.270
<i>Panel B. Percentage change in cross-city Gini of local real GDP</i>		
Both channels	2.387	0.030
Search channel	0.830	−0.270
Leisure channel	1.410	0.235
<i>Panel C. Percentage change in national welfare</i>		
Both channels	3.140	8.146
Search channel	0.554	2.155
Leisure channel	2.574	5.984

Notes: This table reports percentage changes relative to the no-HSR baseline. “Network in 2015” uses the historical HSR network observed in 2015, while “Universal network” connects every city pair using the travel-time construction described in the text. The search-only and leisure-only counterfactuals hold the other travel friction at its no-HSR level. Panel A reports the change in national real wage. Since the total number of labor is fixed, the change in national real wage is the same as the change in national real GDP. Panel B reports the change in the Gini coefficient of city-level local real GDP, as defined in Equation (48). Panel C reports the change in national welfare.

directed city pairs remained unconnected in 2015, consistent with the roles of political, geographic, and other planning considerations.

To isolate the economic implications of a pure volume-based prioritization rule, we construct a sequence of counterfactual networks in which links are added strictly in descending order of baseline business-travel volume. They provide a benchmark for evaluating how rapidly production, welfare, migration, and spatial concentration accumulate when important links are prioritized.

For this counterfactual exercise, we define each potential HSR link as an unordered city pair and rank the 45,451 links by total 2011 UnionPay business travel in both directions:³¹

$$TV_{nj} = V_{nj,2011} + V_{jn,2011}.$$

We construct networks covering the top 0.1, 1, 5, 10, 20, 50, and 100 percent of these links. At each coverage level, we apply the corresponding HSR travel-time shocks to all selected city pairs, while all remaining pairs retain their no-HSR travel times.

Table 13 reports the results. The national real-wage gains are concentrated among links with large baseline business-travel volumes. Connecting only the top 0.1% of links raises the national average real wage by 1.176%, equivalent to 44.3% of the 2.657% gain under universal connectivity. At the top 5%, the real-wage gain reaches 2.403%, or 90.4% of the universal benchmark. Expanding coverage from the top 5% to all city pairs therefore adds only 0.254 percentage points to the national average real wage.

The comparison between the historical and universal networks establishes the endpoints of this

³¹Using leisure-travel ranks obtains very similar results.

shift in the sources of real-wage growth; Table 13 further traces how this shift unfolds as network coverage expands. During the early stages of the rollout, strict prioritization of high-volume links generates spatially uneven gains. These links reduce search costs along the busiest corridors while disproportionately improving the relative attractiveness of favored cities. The resulting reallocation of workers toward cities with higher real wages raises the national average real wage through a spatial-composition effect. As coverage broadens, this reallocation margin gradually weakens: the increase in the migrant stock peaks at 8.028% when the top 1% of links are connected and subsequently falls to 0.185% under universal connectivity, while the increase in the spatial Gini of local real GDP declines from 2.592% to 0.030%.

These patterns indicate that the composition of real-wage gains shifts from concentrated search improvements reinforced by labor reallocation toward more broadly based production-side gains from supplier search. Consequently, once the top 5% of links have been connected, further expansion produces only a small incremental increase in the national average real wage, while attenuating the HSR-induced geographic concentration of economic activity.

National welfare follows the same broad pattern: links with larger baseline travel volumes contribute disproportionately to the aggregate gain, but the marginal benefits of network expansion diminish more gradually than for real wages. Connecting the top 5% of links raises welfare by 5.006%, capturing 61.5% of the universal-network gain, while the top 10% generates 6.014%, or 73.8% of that benchmark.

These results do not by themselves identify a desirable stopping point for network construction. That choice depends on construction costs and on how policymakers weight national output, household welfare, and the geographic distribution of economic activity. A policy focused primarily on national real GDP would place greater weight on the early high-volume links, whereas broader coverage remains more valuable when national welfare or regional balance also enters the policy objective.

Table 13: Aggregate effects of the travel-prioritized counterfactual networks

	Top percentage of links						
	0.1	1	5	10	20	50	100
Real wage/GDP	1.176	2.087	2.403	2.472	2.548	2.624	2.657
Total migrants	5.674	8.028	6.745	4.867	3.157	1.389	0.185
Spatial real GDP Gini	1.700	2.592	2.220	1.771	1.273	0.597	0.030
Welfare	0.627	2.739	5.006	6.014	6.911	7.743	8.146

Notes: This table reports outcomes under networks that connect city pairs in descending order of bilateral business-travel volume. All results are measured relative to the no-HSR baseline, with both the business-search and leisure-accessibility channels active. “Total migrants” is the percentage change in the equilibrium stock of workers residing outside their *hukou* city. Because national population is fixed, the change in the national average real wage is also the change in national real GDP. The percentage change in the cross-city Gini coefficient of local real GDP, as defined in Equation (48), is also reported. National welfare is calculated as 100 times its log change.

8 Broader implications for infrastructure policy

Passenger infrastructure expands the geographic reach of economic activity. Firms can draw on suppliers in other cities while keeping production in place, and households can enjoy distant amenities from their current residence. Temporary travel thereby changes the value of locations as well as the flows between them. The HSR counterfactuals trace how this expanded reach affects production, residential choice, and welfare.

Migration responds to relative differences in accessibility across cities. During the historical rollout, uneven improvements in access made some cities more attractive and induced workers to relocate. Broader network expansion spreads these improvements across more cities, compressing the spatial differences that generate migration. Welfare can continue to rise during this process because residents gain access to more destinations from wherever they live. Population change consequently captures the reallocation induced by a network, while welfare also includes gains realized by households that remain in place.

Passenger connectivity also enters production through a margin distinct from freight transportation. Freight costs determine the expense of shipping inputs after a trading relationship has formed. Passenger travel affects firms' ability to find suppliers, assess their products, and transact with them. The no-search benchmark removes this margin and produces substantially smaller real-income gains. This comparison gives business travel a productive role: it helps determine the sourcing opportunities available to firms and, through them, the organization of goods trade.

The production and leisure channels favor different patterns of network expansion. Commercially important links generate concentrated real-wage gains by connecting large volumes of sourcing activity. A relatively small set of these links captures much of the production benefit associated with wider connectivity. Leisure-accessibility gains accumulate across a broader set of origins and destinations. Later additions to the network may deliver modest production gains while still expanding the consumption opportunities available to households. Network design therefore affects both the magnitude and the composition of the returns to passenger infrastructure.

The relevant exposure also extends beyond whether a city has a station. Travel-time reductions vary substantially across connected cities, as do the economic and recreational opportunities at the other ends of the affected routes. Existing trade and migration linkages transmit these changes to cities outside the network. A city's response reflects its position within the network, the destinations made more accessible, and the equilibrium adjustments taking place elsewhere. Station access alone compresses these different margins into a coarse measure of treatment.

The population results illustrate the importance of these network relationships. Large connected cities gain population, with most of their inflows originating in cities outside the network. Small connected cities experience little net population change, although the aggregate figure conceals offsetting inflows and outflows. HSR thus changes the composition and direction of migration even in places where total population remains stable. The distributional effects follow the intensity and topology of network exposure more closely than a simple division between large and small connected cities.

Infrastructure appraisal must consolidate these responses within a common welfare measure.

Adding changes in trade, wages, land values, population, and travel time as separate benefits risks counting the same underlying gain several times. Some outcomes also reflect capitalization or redistribution across locations. The spatial-equilibrium model traces these adjustments back to the original travel-time shock and provides a consistent measure of their combined welfare consequences.

As a rough benchmark, the annual welfare-equivalent gain under the 2015 network is strikingly close to the cumulative construction expenditure incurred by that year. The ratio is approximately 0.88–1.31.³² This benchmark compares a one-year flow of welfare gains with a cumulative capital outlay. A complete benefit–cost analysis would incorporate the future stream of benefits, route-specific construction and operating costs, financing distortions, depreciation, environmental effects, and adjustment costs. It would also compare the realized network with feasible alternative designs and construction sequences.

The same economic forces arise in other settings where connectivity allows people to reach opportunities located elsewhere. Air-service expansions, conventional passenger rail, intercity roads, and reductions in border frictions can all change firms’ access to potential partners and households’ access to consumption opportunities. Digital communication may replace some trips and facilitate others, particularly when firms identify partners remotely and subsequently meet in person. The spatial organization of economic activity depends on the ease of accessing distant opportunities as well as on where those opportunities are located.

Several margins remain for future work. The UnionPay data identify where and when cards are used, while trip purpose must be inferred from travel dates. The model represents travel-mediated procurement through a particular search technology. Data linking business travel directly to the formation and performance of buyer–supplier relationships would provide a sharper test of this mechanism. On the household side, the model determines leisure destinations conditional on travel and abstracts from trip frequency, duration, and expenditure. The counterfactuals also compare long-run equilibria, leaving the dynamics of supplier formation, migration, and infrastructure construction for future research.

9 Conclusion

This paper studies passenger mobility as a determinant of firms’ sourcing decisions, households’ access to amenities, and the resulting spatial equilibrium. We develop a quantitative model in which firms choose where to search for suppliers and households value leisure destinations outside their city of residence. We estimate the model using China’s HSR expansion and city-pair data on passenger and freight travel times, business and leisure trips, trade, and migration. The estimates show that

³²China’s final-revised nominal GDP in 2015 was RMB 68.9052 trillion, according to the National Bureau of Statistics (https://www.stats.gov.cn/xw/tjxw/tzgg/202302/t20230202_1893850.html, accessed September 2026). Applying the model-implied welfare gain of 3.14% gives an annual welfare-equivalent gain of approximately RMB 2.16 trillion. We approximate cumulative construction expenditure by multiplying the 19,000 kilometers of HSR in operation at the end of 2015, as reported in the 2015 Railway Statistical Bulletin (https://www.mot.gov.cn/shuju/tongjishuju/tielu/202512/t20251227_4191829.html, accessed September 2026), by World Bank estimates of average construction costs of RMB 87 million per kilometer for lines designed for 250 km/h and RMB 129 million per kilometer for lines designed for 350 km/h (Ollivier et al., 2014). This yields an estimated construction cost of RMB 1.65–2.45 trillion.

lower passenger travel time increases business travel and bilateral trade after accounting for freight travel time, while directing leisure trips toward destinations that become easier to reach.

In the counterfactuals, the HSR network in place by 2015 raises national average real wages by approximately 2.185% and welfare by approximately 3.14%. Supplier search contributes materially to the real-wage gain, while improved leisure accessibility accounts for most of the welfare gain. The historical rollout also reallocates population toward large connected cities, with most of the inflows originating in cities outside the network. As coverage expands further, accessibility gains become more evenly distributed and the induced migration and spatial concentration recede even as welfare continues to rise. Network expansion can therefore change the composition of its benefits, shifting from spatial reallocation toward wider access to nonlocal amenities.

Passenger networks reshape economic geography by changing the productive and consumption opportunities available from each location. Their effects depend on the routes created, the destinations connected, and the responses induced throughout the network. Infrastructure appraisal should incorporate the resulting changes in supplier access, leisure access, and relative city attractiveness alongside ridership and direct travel-time savings. The economic value of a passenger network ultimately lies in the opportunities it brings within reach.

References

- Ahlfeldt, G. M. and A. Feddersen (2018). From periphery to core: Measuring agglomeration effects using high-speed rail. *Journal of Economic Geography* 18(2), 355–390.
- Ahlfeldt, G. M., S. J. Redding, D. M. Sturm, and N. Wolf (2015). The Economics of Density: Evidence From the Berlin Wall. *Econometrica* 83(6), 2127–2189.
- Albert, C. and J. Monras (2022). Immigration and Spatial Equilibrium: The Role of Expenditures in the Country of Origin. *American Economic Review* 112(11), 3763–3802.
- Allen, T. (2014). Information Frictions in Trade. *Econometrica* 82(6), 2041–2083.
- Allen, T. and C. Arkolakis (2014). Trade and the Topography of the Spatial Economy. *The Quarterly Journal of Economics* 129(3), 1085–1140.
- Anderson, J. E. and E. van Wincoop (2004). Trade Costs. *Journal of Economic Literature* 42(3), 691–751.
- Antràs, P., S. J. Redding, and E. Rossi-Hansberg (2023). Globalization and Pandemics. *American Economic Review* 113(4), 939–981.
- Arkolakis, C. (2010). Market Penetration Costs and the New Consumers Margin in International Trade. *Journal of Political Economy* 118(6), 1151–1199.
- Bernard, A. B., A. Moxnes, and Y. U. Saito (2019). Production Networks, Geography, and Firm Performance. *Journal of Political Economy* 127(2), 639–688.
- Caliendo, L. and F. Parro (2015). Estimates of the Trade and Welfare Effects of NAFTA. *The Review of Economic Studies* 82(1), 1–44.
- Campante, F. and D. Yanagizawa-Drott (2018). Long-Range Growth: Economic Development in the Global Network of Air Links. *The Quarterly Journal of Economics* 133(3), 1395–1458.
- Che, Z., J. Che, and S. Wang (2024). The gains from changes in internal trade costs: A quantitative analysis of China. *Journal of Regional Science* 64(3), 896–930.
- Combes, P.-P., G. Duranton, and L. Gobillon (2019). The Costs of Agglomeration: House and Land Prices in French Cities. *The Review of Economic Studies* 86(4), 1556–1589.
- Cristea, A. D. (2011). Buyer-seller relationships in international trade: Evidence from U.S. States' exports and business-class travel. *Journal of International Economics* 84(2), 207–220.

- Dekle, R., J. Eaton, and S. Kortum (2007). Unbalanced Trade. *American Economic Review* 97(2), 351–355.
- Diao, M. (2018). Does growth follow the rail? The potential impact of high-speed rail on the economic geography of China. *Transportation Research Part A: Policy and Practice* 113, 279–290.
- Dingel, J. I. and F. Tintelnot (2026). Spatial Economics for Granular Settings. *Econometrica* 94(2), 407–464.
- Donaldson, D. (2018). Railroads of the Raj: Estimating the Impact of Transportation Infrastructure. *American Economic Review* 108(4-5), 899–934.
- Donaldson, D. and R. Hornbeck (2016). Railroads and American Economic Growth: A “Market Access” Approach. *The Quarterly Journal of Economics* 131(2), 799–858.
- Eaton, J. and S. Kortum (1999). International Technology Diffusion: Theory and Measurement. *International Economic Review* 40(3), 537–570.
- Eaton, J. and S. Kortum (2002). Technology, Geography, and Trade. *Econometrica* 70(5), 1741–1779.
- Eaton, J., S. Kortum, and F. Kramarz (2026). Firm-to-Firm Trade: Imports, Exports, and the Labor Market. *Econometrica* 94(4), 1135–1170.
- Epple, D., B. Gordon, and H. Sieg (2010). A New Approach to Estimating the Production Function for Housing. *American Economic Review* 100(3), 905–924.
- Faber, B. (2014). Trade Integration, Market Size, and Industrialization: Evidence from China’s National Trunk Highway System. *The Review of Economic Studies* 81(3), 1046–1070.
- Faber, B. and C. Gaubert (2019). Tourism and Economic Development: Evidence from Mexico’s Coastline. *American Economic Review* 109(6), 2245–2293.
- Fang, H., L. Wang, and Y. Yang (2025). Competition and Quality: Evidence from High-Speed Railways and Airlines. *The Review of Economics and Statistics* 107(2), 494–509.
- Fang, M. and Z. Huang (2022). Migration, housing constraints, and inequality: A quantitative analysis of China. *Labour Economics* 78, 102200.
- Giroud, X. (2013). Proximity and Investment: Evidence from Plant-Level Data. *The Quarterly Journal of Economics* 128(2), 861–915.
- Gokan, T., S. Kichko, and J.-F. Thisse (2019). How do trade and communication costs shape the spatial organization of firms? *Journal of Urban Economics* 113, 103191.
- Gumpert, A., H. Steimer, and M. Antoni (2022). Firm Organization with Multiple Establishments. *The Quarterly Journal of Economics* 137(2), 1091–1138.
- Handbury, J. and D. E. Weinstein (2015). Goods Prices and Availability in Cities. *The Review of Economic Studies* 82(1), 258–296.
- Hu, Y. and M. Xu (2024). High-speed Rail and the Demand for Skill: Evidence from Chinese Exports.
- Huang, Y., T. Hong, X. Chang, and T. Ma (2025). Travel and Regional Development: A Quantitative Analysis of China. *Journal of Regional Science* 65(4), 893–927.
- Huang, Y., S. Lu, X. Yang, and Z. Zhao (2018). Exploring Railway Network Dynamics in China from 2008 to 2017. *ISPRS International Journal of Geo-Information* 7(8), 320.
- Huang, Z., L. Yang, J. Li, T. Zhang, Z. Qu, and Y. Miao (2025). Optimizing the train timetable in a high-speed rail corridor: The implications on departure time, fare cost and seat preference of passengers. *PLOS ONE* 20(6), e0326170.
- Kortum, S. S. (1997). Research, Patenting, and Technological Change. *Econometrica* 65(6), 1389–1419.
- Koster, H. R. A., K. Hayakawa, TABUCHI. Takatoshi, and J.-F. Thisse (2021). High-speed Rail and the Spatial Distribution of Economic Activity: Evidence from Japan’s Shinkansen.
- Koster, H. R. A., T. Tabuchi, and J.-F. Thisse (2022). To be connected or not to be connected? The role of long-haul economies. *Journal of Economic Geography* 22(4), 711–753.

- Koster, H. R. A. and J.-F. Thisse (2026). High-speed rail and the spatial economy. *Regional Science and Urban Economics*, 104192.
- Lee, K. H. and B. J. Tan (2026). Urban Transit Infrastructure and Inequality. *The Review of Economics and Statistics* 108(4), 1067–1084.
- Liang, W., R. Song, and C. Timmins (2024). Frictional Sorting: The Impacts of Dual Constraints on Mobility and Housing Supply in China. *International Economic Review* 65(4), 1747–1776.
- Lin, Y. (2017). Travel costs and urban specialization patterns: Evidence from China’s high speed railway system. *Journal of Urban Economics* 98, 98–123.
- Ma, L. and Y. Tang (2024). The distributional impacts of transportation networks in China. *Journal of International Economics* 148, 103873.
- Ma, X. (2022). *Localized Bargaining: The Political Economy of China’s High-Speed Railway Program*. Oxford: Oxford University Press.
- Miyauchi, Y. (2024). Matching and Agglomeration: Theory and Evidence From Japanese Firm-to-Firm Trade. *Econometrica* 92(6), 1869–1905.
- Miyauchi, Y., K. Nakajima, and S. J. Redding (2025). The Economics of Spatial Mobility: Theory and Evidence Using Smartphone Data. *The Quarterly Journal of Economics* 140(4), 2507–2570.
- Oberfield, E. (2018). A Theory of Input–Output Architecture. *Econometrica* 86(2), 559–589.
- Ollivier, G., J. Sondhi, and N. Zhou (2014). High-Speed Railways in China: A Look at Construction Costs. Technical report, World Bank, Beijing.
- Qin, Y. (2017). ‘No county left behind?’ The distributional impact of high-speed rail upgrades in China. *Journal of Economic Geography* 17(3), 489–520.
- Rao, D. (2022). The Role of Environmental Amenities in the Urban Economy: Evidence From a Spatial General Equilibrium Approach.
- Redding, S. J. and M. A. Turner (2015). Transportation Costs and the Spatial Organization of Economic Activity. In G. Duranton, J. V. Henderson, and W. C. Strange (Eds.), *Handbook of Regional and Urban Economics*, Volume 5 of *Handbook of Regional and Urban Economics*, pp. 1339–1398. Elsevier.
- Song, Y. (2014). What should economists know about the current Chinese hukou system? *China Economic Review* 29, 200–212.
- Startz, M. (2021). The Value of Face-to-Face: Search and Contracting Problems in Nigerian Trade.
- Tombe, T. and X. Zhu (2019). Trade, Migration, and Productivity: A Quantitative Analysis of China. *American Economic Review* 109(5), 1843–1872.
- Wu, J., J. Gyourko, and Y. Deng (2012). Evaluating conditions in major Chinese housing markets. *Regional Science and Urban Economics* 42(3), 531–543.
- Zheng, H., J. Többen, E. Dietzenbacher, D. Moran, J. Meng, D. Wang, and D. Guan (2022). Entropy-based Chinese city-level MRIO table framework. *Economic Systems Research* 34(4), 519–544.

A Mathematical appendix

A.1 The microfoundations of the productivity distribution

According to Kortum (1997) and Eaton and Kortum (1999), suppose that ideas arrive as a Poisson process. Then given a time point, the total number of ideas arriving in a location follows a Poisson distribution. Using our notation, it is

$$\Pr \{M_j^k = n\} = \frac{(T_j^k)^n \exp(-T_j^k)}{n!},$$

where T_j^k is the expected count of ideas, reflecting the absolute advantage of the location. Suppose that the efficiency of each idea is independently drawn from a Pareto distribution. Using our notation, it is

$$\Pr \{z_j^k \leq x\} = 1 - x^{-\theta^P},$$

where shape parameter θ^P inversely measures the dispersion of productivity draws, reflecting the comparative advantage of the location. Taken together, the number of ideas with productivity strictly greater than x , denoted by $N_j^k(x)$, is given by

$$\begin{aligned} \Pr \{N_j^k(x) = t\} &= \sum_{n=t}^{\infty} \Pr \{M_j^k = n\} \binom{n}{t} \left(\Pr \{z_j^k \leq x\}\right)^{n-t} \left(\Pr \{z_j^k > x\}\right)^t \\ &= \frac{(T_j^k x^{-\theta^P})^t \exp(-T_j^k x^{-\theta^P})}{t!}. \end{aligned}$$

This is equivalent to directly assuming that the productivity draws constitute a Poisson point process with an intensity measure (expectation) $T_j^k x^{-\theta^P}$.

A.2 Poisson thinning property

Here we formally prove that when a final goods firm f in city n exerts search efforts to randomly encounter a fraction m_{nj}^k of potential intermediate suppliers in city j , the number of encountered suppliers with productivity strictly greater than x , denoted by $N_{jn}^k(x)$, follows a Poisson distribution with expectation $m_{nj}^k \bar{\lambda}_j^k(x)$.

In our model, the production technology of each city is modeled as a Poisson point process. Specifically, for any variety ω^k , the number of potential intermediate suppliers in city j with productivity higher than x follows a Poisson distribution with expectation $\bar{\lambda}_j^k(x) = T_j^k x^{-\theta^P}$:

$$\Pr \{N_{jn}^k(x) = t\} = \frac{\bar{\lambda}_j^k(x)^t}{t!} e^{-\bar{\lambda}_j^k(x)}.$$

The search process implies that each of these t potential suppliers is independently encountered

by the final goods firm with probability m_{nj}^k , and not encountered with probability $1 - m_{nj}^k$. Therefore, conditional on the total number of potential suppliers being $N_j^k(x) = t$, the number of encountered suppliers $N_{jn}^k(x)$ follows a binomial distribution:

$$\Pr \left\{ N_{jn}^k(x) = t^* \mid N_j^k(x) = t \right\} = \binom{t}{t^*} \left(m_{nj}^k \right)^{t^*} \left(1 - m_{nj}^k \right)^{t-t^*}, \quad t \geq t^*.$$

Therefore, the probability mass function of $N_{jn}^k(x)$ is given by

$$\begin{aligned} \Pr \left\{ N_{jn}^k(x) = t^* \right\} &= \sum_{t=t^*}^{\infty} \Pr \left\{ N_{jn}^k(x) = t^* \mid N_j^k(x) = t \right\} \Pr \left\{ N_j^k(x) = t \right\} \\ &= \frac{\left[m_{nj}^k \bar{\lambda}_j^k(x) \right]^{t^*}}{t^*!} e^{-\bar{\lambda}_j^k(x)} \sum_{t=t^*}^{\infty} \frac{\left[(1 - m_{nj}^k) \bar{\lambda}_j^k(x) \right]^{t-t^*}}{(t-t^*)!}. \end{aligned}$$

Recognizing the infinite sum as the Taylor series expansion for the exponential function, we can simplify the expression to

$$\begin{aligned} \Pr \left\{ N_{jn}^k(x) = t^* \right\} &= \frac{\left[m_{nj}^k \bar{\lambda}_j^k(x) \right]^{t^*}}{t^*!} e^{-\bar{\lambda}_j^k(x)} \cdot e^{(1 - m_{nj}^k) \bar{\lambda}_j^k(x)} \\ &= \frac{\left[m_{nj}^k \bar{\lambda}_j^k(x) \right]^{t^*}}{t^*!} e^{-m_{nj}^k \bar{\lambda}_j^k(x)}, \end{aligned}$$

which is exactly a Poisson distribution with expectation $m_{nj}^k \bar{\lambda}_j^k(x)$.

A.3 The uniqueness and global optimality of search efforts

Here we formally prove that the first-order conditions for the final goods firm's search efforts yield a unique global maximum. For notational convenience, we fix the destination city n and sector k and suppress these subscripts and superscripts. Let $m_j \equiv m_{nj}^k$ denote the search effort in origin city j and $\mathcal{Y}(\mathbf{m}) \equiv \mathcal{Y}_n^k(m_{n1}^k, \dots, m_{nN}^k) \in \mathbb{R}_{++}$, where $\mathbf{m} \equiv [m_1, \dots, m_N]' \in \mathbb{R}_{++}^N$ is the vector of search efforts. Additionally, we define the following notations:

$$\begin{aligned} \alpha &\equiv \frac{\sigma - 1}{\theta^P} > 0, \quad A \equiv \left(\frac{\sigma}{\sigma - 1} \right)^{-\sigma} \frac{\gamma^{1-\sigma}}{\sigma - 1} w_n \tilde{L}_n^k \left(p_n^k \right)^{\sigma-1} > 0, \\ B_j &\equiv T_j^k \left(c_j^k \bar{d}_{jn} \right)^{-\theta^P} > 0, \quad C_j \equiv \kappa \tau_{nj} w_n > 0. \end{aligned}$$

Both of these terms are regarded as given in the optimization problem of individual final goods firms. By definition, we have $\mathcal{Y} \equiv \sum_{j=1}^N B_j m_j$. With these notations, the final goods firm's optimization problem can be written as

$$\max_{\mathbf{m} \in \mathbb{R}_{++}^N} f(\mathbf{m}) = A \mathcal{Y}^\alpha - \sum_{j=1}^N C_j m_j^\kappa.$$

Additionally, we impose the parametric restriction $\kappa \theta^P > \sigma - 1$, which is exactly the parametric

restriction required for the free entry condition to yield a positive mass of final goods firms (i.e., $F_n^k > 0$) and is verified in the parameter estimation section.

Finally, as mentioned in the main text, we assume that the scale parameter κ is sufficiently large such that the optimal search effort always falls within $(0, 1)$ as long as it is strictly positive. Therefore, our proof of uniqueness solely focuses on the uniqueness of a strictly positive search effort vector $\mathbf{m}^* \in \mathbb{R}_{++}^N$ without discussing the bounding constraint.

Uniqueness.—The corresponding first-order condition of this optimization problem is given by

$$\frac{\partial f}{\partial m_j} = A\alpha\gamma^{\alpha-1}B_j - \kappa C_j m_j^{\kappa-1} = 0, \quad \forall j \in \mathcal{N}.$$

Therefore, we have

$$m_j^* = \left[\frac{A\alpha B_j}{\kappa C_j} (\gamma^*)^{\alpha-1} \right]^{\frac{1}{\kappa-1}}, \quad (\text{A.1})$$

which is Equation (15) in the main text. This equation implies that solving the optimal search effort is equivalent to solving for an optimal aggregate scalar γ^* . With Equation (A.1), we further have

$$\gamma^* \equiv \sum_{j=1}^N B_j m_j^* = \left[\sum_{j=1}^N \left(\frac{A\alpha}{\kappa C_j} \right)^{\frac{1}{\kappa-1}} B_j^{\frac{\kappa}{\kappa-1}} \right] \cdot (\gamma^*)^{\frac{\alpha-1}{\kappa-1}}. \quad (\text{A.2})$$

Given that $(\alpha - 1)/(\kappa - 1) < 1$ (since $\kappa\theta^P > \sigma - 1$) and $\sum_{j=1}^N (A\alpha/\kappa C_j)^{1/(\kappa-1)} B_j^{\kappa/(\kappa-1)} > 0$, there exists a unique, strictly positive solution for $\gamma^* \in \mathbb{R}_{++}$ that solves Equation (A.2). Therefore, there exists a unique, strictly positive vector $\mathbf{m}^* \in \mathbb{R}_{++}^N$ that solves the first-order condition.

Global optimality.—We start by proving that the Hessian matrix of $f(\mathbf{m})$ is strictly negative definite at $\mathbf{m}^*$, which ensures that $\mathbf{m}^*$ is a strict local maximum point. The Hessian matrix $\mathbf{H} \equiv [H_{ji}]_{N \times N}$ is given by

$$H_{ji}(\mathbf{m}) \equiv \frac{\partial^2 f}{\partial m_j \partial m_i} = \begin{cases} A\alpha(\alpha - 1)\gamma^{\alpha-2}B_jB_i, & i \neq j, \\ A\alpha(\alpha - 1)\gamma^{\alpha-2}B_j^2 - \kappa(\kappa - 1)C_jm_j^{\kappa-2}, & i = j. \end{cases}$$

According to the first-order condition, we have $A\alpha(\gamma^*)^{\alpha-1}B_j/m_j^* = \kappa C_j(m_j^*)^{\kappa-2}$ at $\mathbf{m}^*$. Substituting this into the Hessian matrix, we get

$$H_{ji}(\mathbf{m}^*) = \begin{cases} A\alpha(\alpha - 1)(\gamma^*)^{\alpha-2}B_jB_i, & i \neq j, \\ A\alpha(\gamma^*)^{\alpha-2} \left[(\alpha - 1)B_j^2 - (\kappa - 1)\gamma^* \frac{B_j}{m_j^*} \right], & i = j. \end{cases}$$

To prove that $\mathbf{H}(\mathbf{m}^*)$ is strictly negative definite, we need to show that $\mathbf{v}'\mathbf{H}(\mathbf{m}^*)\mathbf{v} < 0$ for any

non-zero vector $\mathbf{v} \equiv [v_1, \dots, v_N]' \in \mathbb{R}^N$. By definition, we have

$$\begin{aligned} \mathbf{v}'\mathbf{H}(\mathbf{m}^*)\mathbf{v} &= A\alpha (\gamma^*)^{\alpha-2} \left[(\alpha-1) \sum_{j=1}^N \sum_{i=1}^N v_j B_j B_i v_i - (\kappa-1) \gamma^* \sum_{j=1}^N \frac{B_j}{m_j^*} v_j^2 \right] \\ &= A\alpha (\gamma^*)^{\alpha-2} \left[(\alpha-1) \left(\sum_{j=1}^N v_j B_j \right)^2 - (\kappa-1) \gamma^* \sum_{j=1}^N \frac{B_j}{m_j^*} v_j^2 \right]. \end{aligned}$$

Since $\gamma \equiv \sum_{j=1}^N B_j m_j$, using the Cauchy-Schwarz inequality, we have

$$\left(\sum_{j=1}^N v_j B_j \right)^2 = \left(\sum_{j=1}^N \sqrt{B_j m_j^*} \frac{\sqrt{B_j} v_j}{\sqrt{m_j^*}} \right)^2 \leq \left(\sum_{j=1}^N B_j m_j^* \right) \left(\sum_{j=1}^N \frac{B_j}{m_j^*} v_j^2 \right) = \gamma^* \sum_{j=1}^N \frac{B_j}{m_j^*} v_j^2. \quad (\text{A.3})$$

To bound the term in the bracket of the quadratic form, we consider two cases:

1. If $\alpha \leq 1$, we have $\alpha - 1 \leq 0$, which ensures that the term in the bracket is negative. Therefore, we have $\mathbf{v}'\mathbf{H}(\mathbf{m}^*)\mathbf{v} < 0$.
2. If $\alpha > 1$, using Equation (A.3) and given that $\kappa\theta^P > \sigma - 1$ ensures that $\kappa > \alpha$, we have

$$\mathbf{v}'\mathbf{H}(\mathbf{m}^*)\mathbf{v} \leq A\alpha (\gamma^*)^{\alpha-1} (\alpha - \kappa) \sum_{j=1}^N \frac{B_j}{m_j^*} v_j^2 < 0.$$

In both cases, we strictly have $\mathbf{v}'\mathbf{H}(\mathbf{m}^*)\mathbf{v} < 0$. This proves that $\mathbf{m}^*$ is a strict local maximum point.

Next, we prove that $\mathbf{m}^*$ is the global maximum over the feasible domain $\mathbb{R}_{++}^N$. Since $\alpha < \kappa$, the search cost term dominates for large $\mathbf{m}$, implying that $f(\mathbf{m}) \rightarrow -\infty$ as $\|\mathbf{m}\| \rightarrow \infty$. Consequently, for any $M > 0$, we can always define a sufficiently large radius $R > 0$ (with $R > \|\mathbf{m}^*\|$) such that for all $\mathbf{m} \in \mathbb{R}_+^N$ satisfying $\|\mathbf{m}\| \geq R$, we strictly have $f(\mathbf{m}) < f(\mathbf{m}^*) - M$. Let $\mathcal{D} \equiv \{\mathbf{m} \in \mathbb{R}_+^N \mid \|\mathbf{m}\| \leq R\}$. Since $\mathcal{D}$ is bounded and closed (thus compact), and $f(\mathbf{m})$ is continuous on $\mathcal{D}$, the Weierstrass Theorem guarantees that $f(\mathbf{m})$ must attain a global maximum in $\mathcal{D}$.

Because $f(\mathbf{m}) < f(\mathbf{m}^*) - M < f(\mathbf{m}^*)$ everywhere outside $\mathcal{D}$ (i.e., when $\|\mathbf{m}\| > R$), the global maximum over the entire domain $\mathbb{R}_+^N$ must be identical to the global maximum within $\mathcal{D}$. Furthermore, the continuity of $f(\mathbf{m})$ ensures that any $\mathbf{m}'$ on the outer boundary of $\mathcal{D}$ (i.e., where $\|\mathbf{m}\| = R$) satisfies $f(\mathbf{m}') \leq f(\mathbf{m}^*) - M < f(\mathbf{m}^*)$. Therefore, the global maximum cannot be located on this outer boundary. As a result, the global maximum point must lie either in the strictly positive interior (where $\mathbf{m} \in \mathbb{R}_{++}^N$) or on the coordinate boundaries (where $m_j = 0$ for some $j \in \mathcal{N}$).

Suppose that a boundary point $\mathbf{m}'' \in \mathcal{D}$, where $m_j = 0$ for some but not all $j \in \mathcal{N}$, is the global maximum point. Since $\kappa > 1$, we have $\partial f / \partial m_j = A\alpha \gamma^{\alpha-1} B_j > 0$ at $\mathbf{m}''$. This implies that a marginal increase in m_j from 0 to a small positive value will strictly increase $f(\mathbf{m}'')$, and this new point is still in $\mathcal{D}$. Therefore, $\mathbf{m}''$ cannot be the global maximum point.

For the origin point where all $m_j = 0$, we have $f(\mathbf{0}) = 0$. Using the same derivation as in

Appendix A.4, we can show that

$$f(\mathbf{m}^*) = A (\gamma^*)^\alpha \left(1 - \frac{\alpha}{\kappa}\right) > 0.$$

Since $f(\mathbf{m}^*) > f(\mathbf{0})$, the global maximum point cannot be $\mathbf{0}$. Therefore, the unique global maximum point must be the interior stationary point $\mathbf{m}^*$.

A.4 The total optimized search costs

We prove Equation (19) here. Recall that the firm chooses its search effort vector $\{m_{nj}^k\}_{j \in \mathcal{N}}$ to maximize its net profit. Taking the derivative of the gross profit with respect to m_{nj}^k and equating it to the marginal search cost yields the first-order condition:

$$\left(\frac{\sigma}{\sigma-1}\right)^{-\sigma} \frac{1}{\theta^P} \gamma^{1-\sigma} \frac{w_n \tilde{L}_n^k}{(p_n^k)^{1-\sigma}} \left(\gamma_n^k\right)^{\frac{\sigma-1-\theta^P}{\theta^P}} T_j^k \left(c_j^k \bar{d}_{jn}\right)^{-\theta^P} = \kappa \tau_{nj} w_n \left(m_{nj}^k\right)^{\kappa-1}, \quad \forall j \in \mathcal{N}.$$

From the main text, we have shown that in the symmetric free-entry equilibrium, all active final goods firms set an identical price. Therefore, the sector-specific aggregate price index is given by

$$p_n^k = \frac{\sigma}{\sigma-1} c_n^k \left(F_n^k\right)^{\frac{1}{1-\sigma}}.$$

Substituting the unit cost expression $c_n^k = \gamma(\gamma_n^k)^{-1/\theta^P}$ into this price index, we obtain

$$\left(p_n^k\right)^{1-\sigma} = \left(\frac{\sigma}{\sigma-1}\right)^{1-\sigma} \gamma^{1-\sigma} \left(\gamma_n^k\right)^{-\frac{1-\sigma}{\theta^P}} F_n^k.$$

Then, by substituting this expression back into the first-order condition, we have

$$\frac{\sigma-1}{\sigma \theta^P} \frac{w_n \tilde{L}_n^k}{F_n^k} \frac{1}{\gamma_n^k} T_j^k \left(c_j^k \bar{d}_{jn}\right)^{-\theta^P} = \kappa \tau_{nj} w_n \left(m_{nj}^k\right)^{\kappa-1}, \quad \forall j \in \mathcal{N}.$$

Further, we multiply both sides of the equation by the search effort m_{nj}^k to obtain

$$\frac{\sigma-1}{\sigma \theta^P} \frac{w_n \tilde{L}_n^k}{F_n^k} \frac{1}{\gamma_n^k} \left[T_j^k m_{nj}^k \left(c_j^k \bar{d}_{jn}\right)^{-\theta^P} \right] = \kappa \tau_{nj} w_n \left(m_{nj}^k\right)^\kappa, \quad \forall j \in \mathcal{N},$$

where we recognize the expression of trade share on the left-hand side. Therefore, this equation can be rewritten as

$$\frac{\sigma-1}{\sigma \theta^P} \frac{w_n \tilde{L}_n^k}{F_n^k} \pi_{jn}^k = \kappa \tau_{nj} w_n \left(m_{nj}^k\right)^\kappa, \quad \forall j \in \mathcal{N}.$$

Finally, summing this equation over all possible origin cities $j \in \mathcal{N}$ yields

$$\sum_{j=1}^N \kappa \tau_{nj} w_n \left(m_{nj}^k\right)^\kappa = \frac{\sigma-1}{\kappa \sigma \theta^P} \frac{w_n \tilde{L}_n^k}{F_n^k}.$$

A.5 The counterfactual equilibrium in relative changes and algorithm

A.5.1 The counterfactual equilibrium in relative changes

Here we formally define the equilibrium in relative changes and the full system of non-linear equations. Recall that within the Exact Hat Algebra framework, we focus on the changes of variables relative to the baseline equilibrium. For any variable x , let x' denote its counterfactual value and $\hat{x} \equiv x'/x$ denote its relative change.

Define

$$\mathcal{B} \equiv \left\{ \lambda_{m|n}^T, \lambda_{jn}, \bar{L}_j, \pi_{jn}^k, X_n^k, R_n^k, L_n^k \right\}_{j,n \in N, m \in N \setminus \{n\}, k \in \mathcal{K}}$$

as the set of baseline equilibrium allocations,

$$\mathcal{P} \equiv \left\{ \alpha^T \theta^T, \zeta, \beta, \theta^R, \psi / \theta^T, \theta^P, \kappa, \alpha^P, \sigma \right\} \cup \left\{ \rho_n^k, \gamma_n^k, \gamma_n^{k,s} \right\}_{n \in N, k, s \in \mathcal{K}}$$

as the set of structural parameters, and

$$\mathcal{S}' \equiv \left\{ t'_{nj} \right\}_{j,n \in N}, \quad \mathcal{S} \equiv \left\{ t_{nj} \right\}_{j,n \in N}$$

as the sets of exogenous travel costs in the counterfactual equilibrium and baseline equilibrium, respectively, with their ratio $\hat{t}_{nj} = t'_{nj}/t_{nj}$ representing the exogenous shock induced by the HSR network.

We now define the equilibrium under $\mathcal{S}'$ relative to $\mathcal{S}$ as follows:

Definition A.1. Given $\{\mathcal{B}, \mathcal{P}\}$, define $\{\hat{w}_n, \hat{p}_n^k, \hat{L}_n\}_{n \in N, k \in \mathcal{K}}$ as an equilibrium under $\mathcal{S}'$ relative to $\mathcal{S}$. The equilibrium in relative changes satisfies the following equations:

Expected leisure travel utility:

$$\hat{\Phi}_n^T = \sum_{m \in N \setminus \{n\}} \hat{t}_{nm}^{-\alpha^T \theta^T} \lambda_{m|n}^T. \quad (\text{A.4})$$

Housing rents:

$$\hat{r}_n = \hat{w}_n^{1-\zeta} \hat{L}_n^{1-\zeta}.$$

Consumer price index:

$$\hat{P}_n = \left[\prod_{k=1}^K \left(\hat{p}_n^k \right)^{\rho_n^k} \right]^{\beta} \hat{r}_n^{1-\beta}.$$

Migration share:

$$\hat{\lambda}_{jn} = \frac{\hat{w}_n^{\theta^R} \hat{P}_n^{-\theta^R} (\hat{\Phi}_n^T)^{\psi \theta^R / \theta^T}}{\sum_{n'=1}^N \hat{w}_{n'}^{\theta^R} \hat{P}_{n'}^{-\theta^R} (\hat{\Phi}_{n'}^T)^{\psi \theta^R / \theta^T} \lambda_{jn'}}.$$

Labor distribution:

$$\hat{L}_n = \frac{\sum_{j=1}^N \hat{\lambda}_{jn} \lambda_{jn} \bar{L}_j}{\sum_{j=1}^N \lambda_{jn} \bar{L}_j}.$$

Marginal costs of intermediate good production:

$$\hat{c}_j^k = \hat{w}_j^{\gamma_j^k} \prod_{s=1}^K \left(\hat{p}_j^s \right)^{\gamma_j^{k,s}}.$$

Trade share:

$$\hat{\pi}_{jn}^k = \frac{\left(\hat{c}_j^k \right)^{-\frac{\theta^P \kappa}{\kappa-1}} \hat{t}_{nj}^{\frac{\alpha^P}{1-\kappa}}}{\sum_{j'=1}^N \left(\hat{c}_{j'}^k \right)^{-\frac{\theta^P \kappa}{\kappa-1}} \hat{t}_{nj'}^{\frac{\alpha^P}{1-\kappa}} \pi_{j'n}^k}. \quad (\text{A.5})$$

Goods market clearing conditions:

$$\hat{X}_j^k X_j^k = \rho_j^k \hat{w}_j \hat{L}_j w_j L_j + \sum_{s=1}^K \gamma_j^{s,k} \hat{R}_j^s R_j^s,$$

$$\hat{R}_j^s R_j^s = \sum_{n=1}^N \frac{\sigma-1}{\sigma} \hat{\pi}_{jn}^s \pi_{jn}^s \hat{X}_n^s X_n^s.$$

Labor market clearing condition:

$$\hat{w}_j w_j \hat{L}_j L_j = \sum_{k=1}^K \gamma_j^k \hat{R}_j^k R_j^k + \frac{1}{\sigma} \sum_{k=1}^K \hat{X}_j^k X_j^k.$$

Sector-specific production labor:

$$\hat{L}_n^k = \frac{\hat{R}_n^k}{\hat{w}_n}.$$

Demand-related labor measure:

$$\hat{\tilde{L}}_n^k \tilde{L}_n^k = \rho_n^k \hat{L}_n L_n + \sum_{s=1}^K \frac{\gamma_n^{s,k}}{\gamma_n^s} \hat{L}_n^s L_n^s.$$

Sector-specific final good price index:

$$\hat{p}_n^k = \left(\hat{\tilde{L}}_n^k \right)^{\frac{1}{1-\sigma}} \left[\sum_{j=1}^N \left(\hat{c}_j^k \right)^{-\frac{\theta^P \kappa}{\kappa-1}} \hat{t}_{nj}^{\frac{\alpha^P}{1-\kappa}} \pi_{jn}^k \right]^{\frac{1-\kappa}{\kappa \theta^P}}.$$

A.5.2 The numerical algorithm for counterfactual analysis

We implement the counterfactual analysis in three steps. First, we construct an observed reference equilibrium, denoted by $\mathcal{B}^{\text{obs}}$. Due to differences in data availability across outcomes, this equilibrium combines the spatial distribution of labor and migration shares from the 2010 Population Census, leisure travel shares from the 2011 UnionPay data, and sectoral trade shares and input-output relationships from the 2012 MRIO table. We date this composite observed equilibrium as 2011 and interpret it as characterizing the Chinese spatial economy at the beginning of the 2010s.

For $\mathcal{B}^{\text{obs}}$, we obtain the equilibrium shares and the initial distribution of labor based on *hukou*

directly from the data. These observed objects include $\{\lambda_{m|n}^{T,\text{obs}}, \lambda_{jn}^{\text{obs}}, \bar{L}_j, \pi_{jn}^{k,\text{obs}}\}_{j,n \in \mathcal{N}, m \in \mathcal{N} \setminus \{n\}, k \in \mathcal{K}}$. The remaining variables in the observed reference equilibrium are recovered from the goods and labor market clearing conditions described below. At this stage, “observed” refers to the composite early-2010s allocation and does not yet refer to an economy without HSR. For $\mathcal{P}$, we calibrate or estimate all structural parameters, as detailed in Section 4.

The remaining unknown variables in the observed reference equilibrium $\{X_n^k, R_n^k, L_n^k\}_{n \in \mathcal{N}, k \in \mathcal{K}}$ are recovered by solving the goods and labor market clearing conditions (i.e., Equations (25), (26), and (27) in the main text) as follows. To simplify notation, we suppress the superscript on other variables in the observed reference equilibrium:

$$X_j^k = \rho_j^k w_j L_j + \sum_{s=1}^K \gamma_j^{s,k} R_j^s,$$

$$R_j^s = \sum_{n=1}^N \frac{\sigma - 1}{\sigma} \pi_{jn}^{s,\text{obs}} X_n^s,$$

$$w_n L_n^{\text{obs}} = \sum_{k=1}^K \left(\gamma_n^k R_n^k + \frac{1}{\sigma} X_n^k \right).$$

Notice that given $\{L_n^{\text{obs}}, \pi_{jn}^{k,\text{obs}}\}_{j,n \in \mathcal{N}, k \in \mathcal{K}}$ and $\{\sigma, \rho_n^k, \gamma_n^k, \gamma_n^{s,k}\}_{n \in \mathcal{N}, k, s \in \mathcal{K}}$, the remaining variables to be solved for are $\{X_n^k, R_n^k, w_n\}_{n \in \mathcal{N}, k \in \mathcal{K}}$.¹ These relations form a linear system of $2NK + N$ equations with $2NK + N$ unknowns. By imposing the normalization $\sum_{n=1}^N \sum_{k=1}^K R_n^k = 1$, we can exactly solve this linear system to obtain $\{X_n^k, R_n^k, w_n\}_{n \in \mathcal{N}, k \in \mathcal{K}}$.² Because intermediate firms in sector k of city n allocate a fraction γ_n^k of their total revenue to production labor, we solve for $L_n^k = \gamma_n^k R_n^k / w_n$. Finally, we compute $\tilde{L}_n^k$ according to its definition $\tilde{L}_n^k \equiv \rho_n^k L_n + \sum_{s=1}^K (\gamma_n^{s,k} / \gamma_n^s) L_n^s$.

As highlighted in Section 5, the observed reference equilibrium already incorporates the effects of HSR lines operating in and before 2011. Therefore, our second step is to recover a no-HSR baseline equilibrium, denoted by $\mathcal{B}^0$. Let $\hat{t}_{nj,2011} = t_{nj,2011}^{\text{HSR}} / t_{nj}^{\text{no-HSR}}$, where $t_{nj}^{\text{no-HSR}}$ is the minimum travel time by highway or conventional rail in 2011, subject to $\hat{t}_{nj,2011} = 1$ whenever the city pair is not connected by HSR (including local pairs) or HSR is slower than the no-HSR alternative. To remove the HSR network from the observed reference equilibrium, we apply the inverse travel-time shock $\hat{t}_{nj}^{0/\text{obs}} = 1 / \hat{t}_{nj,2011}$.

Taking $\mathcal{B}^{\text{obs}}$ as the baseline allocation in Definition A.1, we then solve our first counterfactual equilibrium, where the shock is the removal of HSR lines (i.e., $\{\hat{t}_{nj}^{0/\text{obs}}\}_{n,j \in \mathcal{N}}$). This equilibrium later serves as our no-HSR baseline equilibrium.

With all parameters and baseline allocation in hand, solving the counterfactual equilibrium becomes a fixed-point problem over a system of non-linear equations. We explicitly formulate this fixed-point problem as a mapping function that evaluates system residuals. Specifically, our

¹The distribution of labor L_n^{obs} is directly derived as $L_n^{\text{obs}} = \sum_{j=1}^N \lambda_{jn}^{\text{obs}} \bar{L}_j$.

²This normalization is mathematically necessary to secure a unique root because the system of market clearing conditions is homogeneous of degree one in nominal values. We perform a symmetric normalization in the counterfactual equilibrium. Therefore, while real variables are comparable, the absolute nominal values of wages and revenues are not directly comparable across equilibria.

numerical algorithm proceeds as follows: First, we construct an initial guess vector for the relative changes in wages, final goods prices, and labor distributions, denoted as $\mathbf{v}^0 = \{\hat{w}_n, \hat{p}_n^k, \hat{L}_n\}_{n \in \mathcal{N}, k \in \mathcal{K}}$ (in practice, we set all of them to 1). Second, taking $\mathbf{v}^0$ as the given input, we sequentially evaluate the equations in Definition A.1 to compute the relative changes of all other endogenous variables, while strictly imposing the counterfactual numeraire condition $\sum_{n=1}^N \sum_{k=1}^K R_n^k \hat{R}_n^k = 1$. Then, the system yields a newly implied set of values for wages, prices, and labor, which we denote as $\mathbf{v}^1$. Finally, we define the output of our mapping function as the discrepancy between the implied and guessed values: $\mathcal{M}(\mathbf{v}^0) = \mathbf{v}^1 - \mathbf{v}^0$. We deploy a Newton–Krylov root-finding algorithm to iteratively update the guess until the residual norm converges to zero (in practice, we set the tolerance to 1×10^{-6}). Once converged, the terminal vector constitutes the counterfactual equilibrium in relative changes.

For each endogenous variable x , denote the resulting relative change by $\hat{x}^{0/\text{obs}} \equiv x^0/x^{\text{obs}}$. We then recover the corresponding no-HSR equilibrium levels as $x^0 = \hat{x}^{0/\text{obs}} x^{\text{obs}}$. These levels jointly define the no-HSR baseline allocation set $\mathcal{B}^0$. This allocation preserves the structural parameters and fundamentals of the observed early-2010s economy while removing the effects of the HSR network.

Third, we use $\mathcal{B}^0$ as the baseline equilibrium for all annual HSR counterfactuals. For each year $t \in \{2008, \dots, 2015\}$, the travel-time shock relative to the no-HSR baseline is $\hat{t}_{njt} = t_{njt}^{\text{HSR}}/t_{nj}^{\text{no-HSR}}$, again setting the ratio to one for unconnected and local pairs and whenever HSR does not reduce travel time. Using the same numerical procedure as above, we solve Definition A.1 using $\mathcal{B}^0$ as the baseline allocation. The collection of solved relative changes characterizes the long-run spatial equilibrium associated with the HSR network observed in year t , relative to the common no-HSR baseline. Because $\mathcal{B}^0$ is constructed by applying the inverse of the 2011 HSR travel-time shock to $\mathcal{B}^{\text{obs}}$, applying the 2011 HSR shock to $\mathcal{B}^0$ reproduces the observed reference equilibrium up to numerical tolerance.

B Additional data description and construction

B.1 Demographic data

We use gridded demographic data to construct the population centroid for each city, which serves as the city center for calculating pairwise city-to-city travel time. These data are drawn from the Gridded Population of the World (GPWv4) dataset, specifically Version 4.11, provided by NASA. This dataset offers population-count estimates for each grid cell, based on the 2010 national census, at a spatial resolution of $0.0083^\circ \times 0.0083^\circ$, publicly available at <https://www.earthdata.nasa.gov/data/catalog/sedac-ciesin-sedac-gpwv4-apct-wpp-2015-r11-4.11> (accessed December 2025).

The location of each city’s population centroid is calculated as follows:

$$(x_c, y_c) = \frac{\sum_{g \in \mathcal{G}_c} \text{Population}_g \cdot (x_g, y_g)}{\sum_{g \in \mathcal{G}_c} \text{Population}_g},$$

where $\mathcal{G}_c$ represents the set of all grid cells within city c , Population_g denotes the population count in cell $g \in \mathcal{G}_c$, and (x_g, y_g) corresponds to its longitude and latitude coordinates. Consequently, the coordinates of city c ’s population centroid, (x_c, y_c) , are effectively a “center of mass” weighted by population count. Population centroids of 337 cities, including 333 prefecture-level administrative regions and 4 direct-administered municipalities (Beijing, Tianjin, Shanghai, and Chongqing), are presented in Figure B.1.

B.2 Highway network data

To capture the evolution of road connectivity and calculate city-to-city pairwise highway travel time, we utilize geographic data on China’s highway and first-rate road networks from 2008 to 2015, originally constructed by Ma and Tang (2024), which are available at https://github.com/malin84/transportation_networks_of_china (accessed July 2026). This dataset encompasses the complete evolution of highways and first-rate roads from 1994 to 2024.³ The foundational infrastructure data was digitized from comprehensive national and provincial transportation maps.

These maps are initially skeletonized into one-pixel lines and treated as connected graphs, where the physical dimension of each pixel is approximately 0.51 kilometers per edge. Then, the authors divide the network into distinct “segments” bounded by the network’s branch points and endpoints. At the segment level, each segment is assigned a specific road class (highway or first-rate road) and an exact construction year. This temporal vintage determines the specific historical revision of the Technical Standard of Highway Engineering (1988, 1997, 2003, or 2014) applicable to that given segment. Crucially, these engineering standards explicitly dictate the design speeds based on the road class and the underlying terrain.

Moving to the micro-spatial level, each segment is decomposed into a sequence of pixels. At this pixel level, topographic data are overlaid to compute the slope and local elevation range. This

³In our empirical analysis, we classify both highways and first-rate roads under the unified category of the highway network. For expositional simplicity, we refer to them collectively as “highways” throughout the text.




Figure B.1: Location of city population centroids

Notes: This figure illustrates the location of population centroids of 337 cities, including 333 prefecture-level administrative regions and 4 direct-administered municipalities (Beijing, Tianjin, Shanghai, and Chongqing). Each red dot represents the population centroid of a city.

allows the underlying terrain of each pixel to be categorized into plains, low-rolling hills, hills, or mountains. Finally, by intersecting the segment-level engineering standards with the pixel-level terrain categories, a precise design speed is assigned to each pixel. The authors construct an annual dataset documenting the coordinates and speed of each pixel traversed by highways or first-rate roads. For example, in 2015, there are 388,434 pixels identified as part of the highway network. This dataset constitutes the primary raw input for our data construction.

B.2.1 Speed raster construction

We construct the annual speed rasters using a hierarchical overlay approach on a uniform grid matrix. We first establish a high-resolution grid that spans the entire spatial extent of China, defined by its outermost geographic boundaries, with a resolution of $0.005^\circ \times 0.005^\circ$ (approximately 500 meters by 500 meters at the equator). To ensure universal spatial connectivity and prevent isolated nodes, we initialize the entire rectangular bounding box matrix with a universal background speed of 10 km/h, which primarily captures offshore water bodies and unmapped background space. Next, we overlay the Chinese terrestrial polygon onto this matrix, assigning a baseline speed of 30 km/h to all land cells to model the local road network.

We then add the digitized highway network to the matrix by assigning each intersecting grid cell its corresponding engineering design speed. Because our uniform grid and Ma and Tang (2024)’s pixel-level dataset are based on slightly different coordinate reference systems, this spatial mismatch occasionally produces isolated 1-cell gaps along continuous highway corridors. To address this problem, we apply a morphological dilation filter (using a 3×3 local maximum window) to bridge these microscopic gaps and ensure strict topological connectivity. Since our subsequent shortest-path calculations seek time-minimizing routes, this topological correction resolves the network discontinuities without introducing any notable bias to the travel time estimates. Finally, to account for the geographic isolation of Hainan Province, we explicitly code a maritime connection between Xuwen Port and Haikou Port, assigning it an effective speed of 30 km/h.

Following the aforementioned procedure, we construct annual speed raster datasets spanning from 2008 to 2015. Each raster dataset consists of a matrix with dimensions of $10,031 \times 12,318$. As an illustrative example, we present the network graph of the national highway network for the year 2015 in Figure B.2.

B.2.2 Least-cost path algorithm

With the annual speed rasters established, we calculate the shortest travel time between any two city centers using a least-cost path algorithm. The intuition behind this algorithm is to conceptualize the continuous spatial grid as a dense network graph. Each grid cell represents a node, and the transition to any adjacent cell incurs a “time cost” equal to the physical distance divided by the cell’s speed. The algorithm then employs a Dijkstra-type search to find the sequence of adjacent cells that minimizes the cumulative travel time between an origin and a destination.⁴

⁴In practice, we use the Python library `skimage.graph.MCP_Geometric`.




Figure B.2: Highway networks in 2015

Notes: This figure illustrates the 2015 highway network, based on the speed raster data constructed by Ma and Tang (2024). The underlying spatial data comprise a $10,031 \times 12,318$ raster matrix. Grid cells traversed by the highway network are highlighted in red, with the color intensity corresponding to the assigned engineering speed of the respective cell, ranging from 60 km/h to 120 km/h.

Crucially, because our uniform coordinate grid exhibits geographic distortion—where the physical east–west distance of a 0.005° cell shrinks as latitude increases, while the north–south distance remains constant—we dynamically correct the step size for each grid row. The time cost to traverse a cell is calculated using the geometric mean of its latitude-corrected east–west and constant north–south physical distances.

Applying this algorithm, we compute the minimum highway travel time between the city centers of any two cities, which serves as the city-to-city pairwise highway travel time in our analysis.

B.2.3 Passenger versus freight speeds

Since passenger and freight vehicles are subject to different speed regulations, we construct two separate sets of speed rasters. For passenger travel, we preserve the original design speeds. However, according to the *Implementing Regulations of the Road Traffic Safety Law of China*, large motor vehicles (including freight trucks) are strictly prohibited from exceeding 100 km/h on highways.⁵ Accordingly, we construct the freight speed rasters by capping all grid cells with speeds higher than 100 km/h down to the regulatory maximum of 100 km/h. Using the same algorithm, we compute the city-to-city pairwise highway transport time for freight.

To validate our methodology, we compare our estimated city-to-city pairwise highway travel time in 2015 with those calculated by Ma and Tang (2024), as illustrated in Figure B.3. Although our approaches employ slightly different definitions for city centers and imputed travel speeds for off-network pixels, the two sets of estimates exhibit strong consistency. In the subsequent subsection, we detail how we further leverage these highway raster datasets to refine the calculation of railway travel time.

B.3 Railway schedule data

To construct a comprehensive panel of intercity railway travel time, we utilize official train schedule snapshots for each year from 2008 to 2015. The exact snapshot dates are May 5, 2008; July 19, 2009; October 3, 2010; October 28, 2011; December 15, 2012; December 23, 2013; October 26, 2014; and April 9, 2015. Due to historical data availability, the snapshot dates vary across years. However, since our study primarily aims at quantifying the general equilibrium effects during the cumulative construction of China’s HSR rather than capturing the effects of a specific year, this irregularity does not fundamentally undermine our analysis.

The raw data provide sequential timetables for all operational passenger trains in China on the snapshot date, detailing the ordered list of stops, alongside arrival and departure times for each station. Table B.1 shows an example of the train schedule. It is important to note that China’s intercity railway system operates under a rigid, centralized command-and-control framework (Lin, 2017; Qin, 2017). The frequencies, routing, and exact travel times are strictly dictated by the National Railway Train Operating Diagram (*quanguo tielulieche yunxingtu*). Consequently, between official and periodic national schedule adjustments, the daily operations of regular passenger trains are

⁵Available at https://www.gov.cn/zwqk/2005-05/23/content_214.htm (accessed July 2026).




Figure B.3: Comparison of pairwise highway travel time between our calculation and Ma and Tang's in 2015

Notes: This figure compares the city-to-city highway travel times in 2015 computed using our methodology with those estimated by Ma and Tang (2024). Each point represents a unique, undirected city pair. Given that bilateral travel times are symmetric by construction, we restrict the plot to the lower triangular elements of the travel time matrix.

essentially deterministic and fully fixed (Huang et al., 2025). Therefore, a single-day snapshot reliably represents the stable transport connectivity for that period.

Table B.1: An example of the train schedule: D123 on May 5, 2008

Train number	Order	Station name	Arrival time	Departure time
D123	1	Beijing West	09:47	09:47
D123	2	Baoding	10:45	10:46
D123	3	Shijiazhuang	11:47	11:49
D123	4	Anyang	13:22	13:23
D123	5	Luohe	15:37	15:38
D123	6	Xinyang	16:38	16:39
D123	7	Hankou	18:30	18:30

Notes: This table shows an example of the train schedule: D123 on May 5, 2008. The order of the stations is the sequence of the stations the train stops at. The arrival time and departure time are the time of the train’s arrival and departure at the station, respectively. Typically, the train stops at a station for a few minutes.

Converting these operational schedules into a reliable city-to-city origin–destination (OD) travel time matrix requires two major data processing steps: (1) matching each railway station to its corresponding prefecture-level city and geographic coordinates, and (2) transforming sequential timetables into pairwise travel times.

B.3.1 Station-to-city mapping and geocoding

The raw schedule data contain 4,045 unique station names but lack explicit geographic identifiers. To map these stations to their respective cities, we systematically combine information from specialized railway websites (e.g., <https://www.gaotie.com.cn/huochezhan/>) and extract city information for each station from Baidu Baike (a comprehensive online encyclopedia in China that maintains entries for most railway stations, <https://baike.baidu.com/>). We cross-validate the city information retrieved from these sources, manually resolve any inconsistencies, and manually identify the cities for any stations missing from all sources. Finally, we match 4,024 of them with their respective cities. The remaining 21 stations are not located in any prefecture-level administrative region or in one of the four direct-administered municipalities.

Subsequently, we combine the city information and station names to query the Baidu Maps API (<https://api.map.baidu.com/>) to obtain precise longitude and latitude coordinates for each station. To ensure rigorous spatial accuracy, we validate these geocoded coordinates against the official prefecture-level administrative boundaries (see Appendix B.7 for data source) by performing a spatial join to verify whether a station’s coordinates fall strictly within the polygon of its assigned city. A small fraction of stations (approximately 1.8%) exhibiting geographic inconsistencies in this validation step are flagged as potential mismatches and excluded from the subsequent analysis. Figure B.4 shows the location of these identified railway stations in our dataset.

B.3.2 Constructing the station-to-station OD dataset

Our next step is to convert the sequential timetables into pairwise travel times. We first develop an iterative algorithm to infer the absolute time elapsed since the train’s initial departure. We set that




Figure B.4: Location of railway stations

Notes: This figure illustrates the location of 3,953 railway stations in China, which are identified in our railway schedule snapshot data from 2008 to 2015. Each red dot represents a railway station.

every train departs from its initial station on Day 1. Moving sequentially along the train's route, we track the time of day in minutes. Whenever a recorded arrival or departure time is strictly earlier in the day than the preceding recorded time, we infer that the train has crossed midnight. Using this inferred day counter, we convert all times into absolute minutes elapsed since 00:00 of Day 1.

After the conversion of time, we expand the sequential route of each train into all possible pairwise combinations of stations (where the origin precedes the destination). For example, in the example of Table B.1, the expanded pairwise combinations of stations are: (Beijing West, Baoding), (Beijing West, Shijiazhuang), (Beijing West, Anyang), (Beijing West, Luohe), (Beijing West, Xinyang), (Beijing West, Hankou), (Baoding, Shijiazhuang), (Baoding, Anyang), (Baoding, Luohe), (Baoding, Xinyang), (Baoding, Hankou), (Shijiazhuang, Anyang), (Shijiazhuang, Luohe), (Shijiazhuang, Xinyang), (Shijiazhuang, Hankou), (Anyang, Luohe), (Anyang, Xinyang), (Anyang, Hankou), (Luohe, Xinyang), (Luohe, Hankou), (Xinyang, Hankou). For each pairwise combination of stations, we calculate the travel time as the difference between the arrival time of the destination station and the departure time of the origin station.

To ensure the robustness of the travel time network, we apply two strict anomaly detection criteria to filter out erroneous records (e.g., a small number of stations share the same name but are located in different cities):

1. If the stations immediately preceding and succeeding station s on a train's route are both located in city c , but station s itself is assigned to a different city $c' \neq c$, we flag the record at

station s as a routing or mapping anomaly. OD pairs involving such anomalies are dropped.

2. Using the geocoded coordinates, we calculate the minimum distance between the origin and destination stations for every OD pair. We then compute the implied average travel speed by dividing the distance by the travel time. We strictly exclude records implying average speeds that exceed the physical and regulatory limits of the railway system. Specifically, we set a threshold of > 400 km/h for high-speed rail and > 200 km/h for conventional trains (we define the classification of high-speed rail and conventional trains later).

Finally, we aggregate the station-to-station OD data to the city-pair level. Given that multiple train schedules typically operate between any connected city pair, our next step is to identify the minimum rail travel time for each directed pair, consistent with our methodology for computing highway travel times.

B.3.3 Minimum travel time by HSR and conventional trains

To distinguish HSR from conventional trains, we categorize train services based on their official alphanumeric prefixes. Specifically, trains with prefixes G (*Gaotie*, High-Speed EMU), D (*Dongche*, EMU), and C (*Chengji*, Intercity EMU) are defined as HSR, while all other prefixes (e.g., Z, T, K, and numeric-only trains) are classified as conventional trains. For instance, the train detailed in Table B.1 operates with the prefix D and is thus classified as HSR.

This binary classification aligns well with the official operational classifications and general public perception, as G, D, and C trains exclusively utilize Electric Multiple Unit (EMU) technology and offer significantly higher operational speeds (Diao, 2018; Huang et al., 2018). More importantly, it fundamentally captures a structural shift in China's transportation network. For the purpose of our spatial equilibrium analysis, the crucial distinction extends beyond mere travel speed; it reflects the transition from a traditional mixed-use network to a predominantly passenger-dedicated system. Historically, China's conventional railway network functioned as a mixed-use system accommodating both freight and passenger traffic. In contrast, HSR services (G, D, C) operate largely on newly constructed passenger-dedicated lines or upgraded high-speed tracks. By substantially separating passenger traffic from traditional freight networks, this classification accurately maps onto the specific economic channels we parameterize in our structural model: trade search and leisure travel, both of which are fundamentally passenger-driven forms of mobility. Based on this criterion, we split our OD dataset into two mode-specific subsets: one for HSR and one for conventional trains.

After classifying the network, we calculate the minimum travel time for each city pair. Because HSR stations are frequently located in urban peripheries rather than city centers, comparing pure in-vehicle rail time would underestimate the travel time. To ensure comparability across modes, we incorporate the local commuting time to and from the stations. For any given city pair, we identify all direct train schedules. For each available schedule, we define the total travel time as the sum of three components: (1) the road travel time from the origin city center to the departure station, (2) the in-vehicle rail travel time, and (3) the road travel time from the arrival station to the destination city center. The road travel times are computed using the road network rasters constructed in the

previous subsection. Finally, we take the minimum total travel time across all direct schedules as the travel time for that city pair. We repeat this procedure separately to construct the OD travel time matrices for both HSR and conventional trains.

B.4 Migration data

To construct the bilateral labor migration matrix and measure spatial labor mobility, we utilize individual-level microdata from the 2010 National Population Census and the 2015 National 1% Population Sampling Survey. The 2010 dataset provides a representative random sample containing 1,267,381 individuals, while the 2015 dataset provides a similar representative sample covering 1,371,252 individuals. These datasets consistently report respondents' precise current residential location and their official household registration location (*hukou*).

Consistent with our spatial model where individuals belong to an origin city and subsequently choose a destination city to reside and work, we define an individual's origin city j by their *hukou* location, and their destination city n by their current residential location. This aligns with the method used by many studies about China's internal migration (e.g., Tombe and Zhu, 2019; Liang et al., 2024).

With the cities mapped, we classify individuals as intercity migrants if their current residential city differs from their *hukou* city ($n \neq j$). Conversely, individuals whose current residential city coincides with their *hukou* city ($n = j$) are classified as locals. For both the 2010 and 2015 samples, we aggregate the total number of migrants from every origin city j to every destination city n , denoted as L_{jn} , and the total number of locals in each city j , denoted as L_{jj} .

Finally, we construct the directional migration share λ_{jn} . For a given origin city j , we first compute the total number of individuals whose *hukou* originates from j , $\bar{L}_j = \sum_{n=1}^N L_{jn}$. The migration share from city j to city n is then calculated as the fraction of j 's *hukou* population that chooses to reside in n :

$$\lambda_{jn} = \frac{L_{jn}}{\bar{L}_j}.$$

B.5 Intercity, multi-sector input-output table

To empirically capture the spatial production networks and calibrate the structural trade elasticity parameters, we utilize the China City-Level Multi-Regional Input-Output (MRIO) Tables for the years 2012 and 2015, constructed by Zheng et al. (2022) and available at <https://www.ceads.net> (accessed April 2026). The authors construct input-output tables covering the trade flows across 42 economic sectors among 313 regions. Specifically, these tables contain the input values of any sector in any region from any sector in any region, which constitutes a $13,146 \times 13,146$ matrix. Additionally, consistent with the standard structure of input-output table, it also contains the final consumption of any region from any sector in any region (including urban and rural household consumption) and value added (including labor compensation).

It should be noted that the 313 regions in the initial dataset do not belong to a uniform administrative tier. Likely due to data limitations, Hainan, Yunnan, Tibet, and Qinghai Provinces

are not disaggregated into prefecture-level regions, rendering it infeasible to accurately map the HSR network at the desired spatial resolution. Consequently, we exclude these four provinces from our analysis. We further refine the sample by dropping non-prefecture-level regions in the table, such as several provincially administered county-level divisions (*shengzhixixian*). Ultimately, our final dataset comprises 302 cities, consisting of 298 prefecture-level administrative regions and 4 direct-administered municipalities. This set of locations defines the primary spatial domain for our counterfactual analysis.

The 42 sectors comprise 1 primary, 27 secondary, and 14 tertiary sectors. While we retain the disaggregated 42-sector structure for parameter estimation, we aggregate these into the 3 broad macro-sectors. Consequently, the raw trade flow matrix is restructured into a 906×906 matrix for our counterfactual analysis. This preserves the main input–output linkages between agriculture, manufacturing, and services, while keeping the high-dimensional relative-change system numerically tractable.

We detail the data construction procedure as follows. First, we construct the sector-specific bilateral trade share matrix, which maps directly into the intermediate goods trade probability π_{jn}^k . For a specific intermediate good produced by sector k in origin city j and shipped to destination city n , we sum the input values demanded by all utilizing sectors $s \in \mathcal{K}$ within city n . Normalizing this sum by the total expenditures of city n on intermediate goods from sector k across all possible origin cities yields the precise expenditure share, π_{jn}^k . We perform this calculation separately for 2012 and 2015.

Second, we systematically calibrate the underlying structural parameters that dictate the demand (i.e., the expenditure shares and production elasticities). For final consumption, we aggregate the total consumption demands from both rural and urban households in city n across all locations for sector k 's goods. The share of this sector-specific consumption relative to the city's total final consumption yields the Cobb–Douglas preference parameter ρ_n^k .

Finally, we utilize the value-added block of the MRIO tables to calibrate parameters in the intermediate production function. We identify total labor compensation (total wage) for each sector k in city n . By combining the total wage with the intermediate goods input matrix, we calculate the total production input value for each city-sector pair. This allows us to perfectly calibrate the Cobb–Douglas production elasticities: $\gamma_n^{k,s}$ (the cost share of sector s 's intermediate goods in sector k 's production) and γ_n^k (the cost share of labor).

B.6 UnionPay data

We construct bilateral measures of intercity business and leisure travel using anonymized credit- and debit-card transaction records from UnionPay for 2011–2015. During this period, UnionPay operated the nationwide interbank payment network through which cards issued by different Chinese banks were processed. Unlike household surveys or aggregate passenger statistics, the data allow us to follow the same card across locations and dates and thereby construct directional travel flows between prefecture-level cities.

We construct the card sample in two steps. First, we randomly draw 1% of all cards active in

January 2011. In every subsequent month, we supplement this initial cohort with a 1% random sample of cards that newly appear in the transaction records. This sampling procedure yields approximately 27 million distinct cards. For every sampled card, we retrieve its complete transaction history over the study period. Adding newly appearing cards prevents the sample from becoming restricted to the stock of cards already active at the beginning of 2011, while retaining previously sampled cards allows us to follow their locations consistently over time.

The raw records report the date, transaction type, and merchant location associated with each purchase. We retain only offline transactions, which require the cardholder to be physically present at the merchant, and exclude online transactions that do not reveal the cardholder's location. Offline transactions account for approximately 97% of transactions during the relevant period. We map merchant locations to prefecture-level cities. The mobility measures use the occurrence of a transaction rather than its expenditure value, so that travelers making larger purchases do not receive greater weight.

For each card and calendar month m , we assign a single home city based on the card's transaction history within a rolling twelve-month window. Specifically, the home city is the city in which the card is used most frequently from month $m - 6$ through month $m + 5$, comprising the six preceding months, the current month, and the five following months. The same home-city assignment is applied to every transaction made by the card during month m . When the full twelve-month window extends beyond the months available in the transaction records, we use the available months within the corresponding window, so that home assignments near the sample boundaries are based on shorter observation periods. This monthly rolling procedure smooths temporary fluctuations associated with travel while allowing a card's assigned home to change when its user relocates. Conditional on this assignment, a card is recorded as traveling from home city n to destination city j on date d in month m if it makes at least one offline transaction in $j \neq n$ on that date. Multiple transactions by the same card in the same destination on the same date are collapsed into a single card–destination–date observation. Thus, our travel measures count observed traveling cards rather than transactions or expenditure.

The transaction records reveal when and where a card is used but do not report the purpose of the trip. We therefore use calendar timing to construct purpose-specific samples. Trips occurring on non-holiday weekdays are classified as business travel. Trips occurring during national public holidays or an extended Chinese New Year window are classified as leisure travel. We deliberately rely on these relatively clean but non-exhaustive timing windows rather than assigning every observed trip to one of the two purposes.

For each year, we aggregate the resulting card–destination–date observations to the directed city-pair level. The business-travel measure gives the observed bilateral volume V_{nj} used in Equation (36). The timing classification does not imply that every such trip is undertaken for supplier search. As discussed in Section 4.2.3, observed V_{nj} combines procurement-related travel represented by the model with residual business travel, V_{nj}^U , such as administrative visits, conferences, and other corporate trips.

For leisure travel, we normalize the annual flow from origin n to destination m by the total

observed intercity leisure flow originating in n . The resulting destination share is

$$\lambda_{m|n,t}^T = \frac{V_{nm,t}^T}{\sum_{m' \in \mathcal{N} \setminus \{n\}} V_{nm',t}^T}, \quad m \neq n,$$

where $V_{nm,t}^T$ denotes the observed leisure-travel volume from n to m in year t . This measure is the empirical counterpart of the destination-choice probability in Equation (3). We use it to estimate the composite travel-friction elasticity $\alpha^T \theta^T$ and to recover relative destination attractiveness and origin-specific leisure accessibility.

Because the data track cards rather than individuals, they do not cover travelers without payment cards and may count an individual more than once if that person uses multiple sampled cards. They also reveal a traveler's location only when the traveler makes an offline purchase rather than continuously throughout a trip. We therefore interpret the constructed variables as measures of relative bilateral travel flows and destination shares, rather than as exact counts of travelers. Our empirical specifications rely primarily on these relative flows, their changes over time, and comparisons across city pairs constructed using a consistent sampling procedure.

B.7 Other data sources

We utilize prefecture-level average nominal wage data in 2010 and 2015, sourced from the *China City Statistical Yearbook*. Since prefecture-level price index data are not available, we match each prefecture-level nominal wage with consumer price index for its province to construct the real wage. The price index dataset spans from 2000 to 2015, sourced from the *China Statistical Yearbook*. We assume that the price level is the same across all provinces in 1999 and calculate the cumulative price index for 2010, 2011, and 2015.

Our prefecture-level Geographic Information System (GIS) polygon map is obtained from the National Platform for Common GeoSpatial Information Services, administered by China's Ministry of Natural Resources, and is accessible via <https://cloudcenter.tianditu.cn/administrativeDivision> (accessed December 2025).

C Appendix for the estimation of structural parameters

C.1 Sensitivity to the calibration of the migration shape parameter

Our baseline analysis calibrates the migration elasticity to $\theta^R = 1.9$, following the city-level estimate of 1.9 in Fang and Huang (2022). This calibration affects the estimation of the utility weight of leisure travel accessibility through two distinct margins. First, θ^R enters the transformed dependent variable of the migration gravity equation. Second, the estimated coefficient on leisure travel accessibility identifies the composite parameter $\theta^R \psi / \theta^T$, which must be divided by θ^R to recover the structural parameter ψ / θ^T . We therefore examine the sensitivity of our estimate to alternative values of θ^R .

The relevant value of the Fréchet shape parameter depends on the spatial choice margin and the definition of location alternatives. A larger θ^R corresponds to less dispersion in idiosyncratic location preferences and therefore greater substitutability across locations. Locations defined at a finer spatial resolution are typically more similar and closer substitutes than larger geographical units. Consistent with this interpretation, Tombe and Zhu (2019) estimate a migration elasticity of approximately 1.5 across province pairs in China, whereas Fang and Huang (2022) obtain an estimate of 1.9 across China's cities. At a substantially finer spatial scale, Ahlfeldt et al. (2015) obtain a Fréchet shape parameter of 6.83 for commuting decisions across blocks within Berlin.

These estimates should not be interpreted as alternative calibrations of a single parameter that is invariant across spatial settings. In particular, within-city commuting decisions differ substantially from long-run migration across cities in terms of both the relevant alternatives and the underlying mobility frictions. Consequently, we view the city-level estimate of 1.9 as the most appropriate benchmark for our intercity migration setting.

To illustrate analytically how this calibration affects our estimate, consider the regression specification in the main text (i.e., Equation (32)):

$$\ln \lambda_{jnt} - \theta^R \ln \frac{w_{nt}}{P_{nt}} = \theta^R \frac{\psi}{\theta^T} \ln \frac{\Phi_{nt}^T}{\bar{\Phi}_t^T} + \beta_t \ln (\text{Distance}_{jn} + 1) + \delta_{jt} + \epsilon_{jnt},$$

where $\ln \tilde{\varphi}_{nt}^T$ serves as the instrument for $\ln (\Phi_{nt}^T / \bar{\Phi}_t^T)$. To simplify notation, throughout the following derivation, we use the original notation for each variable to denote its residual after partialling out the year-specific distance controls and origin-year fixed effects. The Frisch–Waugh–Lovell theorem implies that the just-identified IV estimator can be written as

$$\begin{aligned} \frac{\overline{\theta^R \psi}}{\theta^T} &= \frac{\sum_{j,n,t} \ln \tilde{\varphi}_{nt}^T [\ln \lambda_{jnt} - \theta^R \ln (w_{nt}/P_{nt})]}{\sum_{j,n,t} \ln \tilde{\varphi}_{nt}^T \ln (\Phi_{nt}^T / \bar{\Phi}_t^T)} \\ &= \frac{\sum_{j,n,t} \ln \tilde{\varphi}_{nt}^T \ln \lambda_{jnt}}{\sum_{j,n,t} \ln \tilde{\varphi}_{nt}^T \ln (\Phi_{nt}^T / \bar{\Phi}_t^T)} - \theta^R \frac{\sum_{j,n,t} \ln \tilde{\varphi}_{nt}^T \ln (w_{nt}/P_{nt})}{\sum_{j,n,t} \ln \tilde{\varphi}_{nt}^T \ln (\Phi_{nt}^T / \bar{\Phi}_t^T)}. \end{aligned}$$

This decomposition separates the migration response associated with instrument-induced variation in leisure accessibility into two components. The first term captures the total migration response per unit of accessibility variation. Because leisure accessibility and real wages can co-move in

spatial equilibrium, this response includes both the direct utility effect of leisure accessibility and the migration response attributable to the associated real-wage differences. The second term removes the latter component by converting the instrument-related real-wage variation into migration responses using the calibrated migration elasticity θ^R . Provided that the instrument is orthogonal to the error term ϵ_{jnt} and that θ^R is correctly calibrated, the remainder identifies the direct leisure-accessibility coefficient $\theta^R \psi / \theta^T$, holding real wages fixed. A value of θ^R that is too large leads to over-correction, whereas a value that is too small leaves part of the real-wage channel attributed to leisure accessibility.

Accordingly, the estimate for ψ / θ^T is given by

$$\widehat{\psi / \theta^T} = \frac{1}{\theta^R} \frac{\sum_{j,n,t} \ln \tilde{\varphi}_{nt}^T \ln \lambda_{jnt}}{\sum_{j,n,t} \ln \tilde{\varphi}_{nt}^T \ln (\bar{\Phi}_{nt}^T / \bar{\Phi}_t^T)} - \frac{\sum_{j,n,t} \ln \tilde{\varphi}_{nt}^T \ln (w_{nt} / P_{nt})}{\sum_{j,n,t} \ln \tilde{\varphi}_{nt}^T \ln (\bar{\Phi}_{nt}^T / \bar{\Phi}_t^T)}.$$

Both ratios on the right-hand side are positive in our data. Therefore, both the estimated coefficient $\widehat{\theta^R \psi / \theta^T}$ and the implied structural parameter $\widehat{\psi / \theta^T}$ decline as the calibrated migration elasticity increases. Importantly, the decline in $\widehat{\psi / \theta^T}$ is not driven solely by mechanically dividing the estimated coefficient by a larger value of θ^R . As discussed earlier, when migration is assumed to respond more strongly to utility differences (i.e., a larger calibrated θ^R), a given migration response implies a smaller underlying utility contribution from leisure accessibility.

Table C.1 re-estimates the pooled IV specification using $\theta^R \in \{1.5, 1.9, 2, 2.5, 3\}$. The implied estimate of ψ / θ^T declines monotonically with θ^R . It remains positive and statistically significant under the province-level calibration of 1.5, the city-level calibration of 1.9, 2, and 2.5. When θ^R is increased further to 3, however, the point estimate becomes smaller and is no longer statistically distinguishable from zero. We interpret the latter value as an upper-end stress test rather than an equally plausible benchmark, given that the most directly comparable city-level estimate is 1.9.

To assess the extent to which our quantitative conclusions depend on the direct utility value of leisure accessibility, Section 5 reports an additional channel-off counterfactual in which we impose $\psi / \theta^T = 0$. This conservative restriction lies below the point estimates obtained under all reported calibrations of θ^R . It removes leisure accessibility from the location-choice equation, so that HSR-induced improvements in leisure accessibility no longer affect migration directly through residential utility. However, migration can continue to respond to changes in real wages generated by the production-side search channel.

Comparing this specification with the full model isolates the quantitative effects of HSR that do not rely on the estimated direct utility value of leisure accessibility. Turning off the direct leisure-accessibility channel attenuates spatial labor reallocation, while reductions in business-travel frictions continue to affect search, trade, prices, real wages, and migration through changes in economic opportunities.

C.2 First-difference specification for migration gravity

As a robustness check to the cross-sectional and pooled specifications in the main text, we also estimate the migration gravity equation in time differences between the two census years. Starting from the structural migration share in Equation (5), the log migration equation in year t can be

Table C.1: Re-estimating the migration gravity equation with different calibrations of θ^R

	Dependent variable: $\ln \lambda_{jnt} - \theta^R \ln \frac{w_{nt}}{P_{nt}}$				
	$\theta^R = 1.5$	$\theta^R = 1.9$	$\theta^R = 2$	$\theta^R = 2.5$	$\theta^R = 3$
	(1)	(2)	(3)	(4)	(5)
$\ln(\widehat{\Phi}_{nt}^T / \widehat{\Phi}_t^T)$	0.209*** (0.060)	0.180*** (0.060)	0.173*** (0.060)	0.137** (0.060)	0.101 (0.063)
$\ln(\text{distance}) \times \mathbb{1}\{t = 2010\}$	-0.895*** (0.017)	-0.916*** (0.017)	-0.922*** (0.018)	-0.948*** (0.019)	-0.974*** (0.021)
$\ln(\text{distance}) \times \mathbb{1}\{t = 2015\}$	-0.858*** (0.016)	-0.874*** (0.016)	-0.878*** (0.016)	-0.897*** (0.017)	-0.917*** (0.018)
Origin $\times$ year fixed effects	Yes	Yes	Yes	Yes	Yes
Estimator	IV	IV	IV	IV	IV
First-stage F-statistic	1,594	1,594	1,594	1,594	1,594
Implied estimate of ψ / θ^T	0.139	0.095	0.087	0.055	0.034
# of observations	28,443	28,443	28,443	28,443	28,443

Notes: This table re-estimates Equation (32) using different calibrations of θ^R . The estimation data are the pooled sample over 2010 and 2015. IV estimator is employed in all columns and Kleibergen–Paap F-statistic is reported. Column (2) replicates the estimation result in Column (6) of Table 4. The implied estimate of ψ / θ^T in each column is calculated by dividing the point estimate of the accessibility term by θ^R and reported in the table. Standard errors, two-way clustered at the origin and destination city levels, are reported in parentheses.

*** p < 0.01, ** p < 0.05, * p < 0.1.

written as

$$\ln \lambda_{jnt} = \theta^R \ln \frac{w_{nt}}{P_{nt}} + \theta^R \frac{\psi}{\theta^T} \ln \Phi_{nt}^T - \theta^R \ln \iota_{jnt} + \theta^R \ln a_{nt} - \ln \Theta_{jt}. \quad (\text{C.1})$$

Differencing Equation (C.1) between $t = 2015$ and $t = 2010$ eliminates time-invariant bilateral migration frictions and yields

$$\Delta \ln \lambda_{jn} - \theta^R \Delta \ln \frac{w_n}{P_n} = \theta^R \frac{\psi}{\theta^T} \Delta \ln \Phi_n^T + \Delta \delta_j + \Delta \epsilon_{jn}, \quad (\text{C.2})$$

where $\Delta x \equiv x_{2015} - x_{2010}$ for any variable x , the origin fixed effect $\Delta \delta_j$ absorbs the change in the multilateral resistance term, and the composite residual $\Delta \epsilon_{jn} \equiv \theta^R \Delta \ln a_n - \theta^R \Delta \ln \iota_{jn}$ collects changes in local residential amenities and migration costs. As in the main text, we impose the calibrated value $\theta^R = 1.9$ and move the real-wage term to the left-hand side, thereby avoiding the need to instrument for endogenous wage changes. Relative to the levels specification, the first-difference regression identifies $\theta^R \psi / \theta^T$ from within-city changes in leisure travel accessibility, rather than from cross-sectional differences in accessibility levels.

C.2.1 Construction of the accessibility difference

Because UnionPay leisure travel data begin in 2011, we again use 2011 measures of leisure accessibility and destination attractiveness as proxies for their 2010 counterparts when constructing $\Delta \ln \Phi_n^T \equiv \ln \hat{\Phi}_n^T$. However, leisure travel accessibility is not directly observed, thus we must

construct $\Delta \ln \Phi_n^T$ based on the estimated leisure travel gravity equation. Recall that

$$\Phi_{nt}^T = \sum_{m \in \mathcal{N} \setminus \{n\}} \left(\frac{b_{mt}}{\mu_{nmt}} \right)^{\theta^T} = \left(\bar{a}^T \right)^{-\theta^T} \sum_{m \in \mathcal{N} \setminus \{n\}} b_{mt}^{\theta^T} t_{nmt}^{-\alpha^T \theta^T}.$$

Therefore, the relative change in accessibility between the two periods satisfies

$$\hat{\Phi}_n^T = \sum_{m \in \mathcal{N} \setminus \{n\}} \hat{b}_m^{\theta^T} \hat{t}_{nm}^{-\alpha^T \theta^T} \lambda_{m|n}^T, \quad (\text{C.3})$$

where $\lambda_{m|n}^T$ denotes the leisure travel share in the earlier period. It should be noted that Equation (C.3) is different from the one in Definition A.1 where $\hat{b}_m = 1$ since HSR does not affect the destination attractiveness in the counterfactual equilibrium. In the regression, $\hat{\Phi}_n^T$ refers to the *observed* relative accessibility between two years, which can be driven by both changes in destination attractiveness and changes in bilateral travel costs.

As shown in Equation (C.3), constructing $\hat{\Phi}_n^T$ requires $\{\lambda_{m|n}^T\}_{n \in \mathcal{N}, m \in \mathcal{N} \setminus \{n\}}$ in 2011. A practical concern is that the observed UnionPay destination shares $\lambda_{m|n}^T$ are measured with substantial sampling noise, especially for thin city pairs. Following the “covariates-based approach” highlighted by Dingel and Tintelnot (2026), we therefore replace the raw shares with their structural counterparts implied by the estimated leisure travel gravity equation. Specifically, given the recovered relative attractiveness $\widehat{b_{mt}^{\theta^T} / \bar{b}_t^{\theta^T}}$ in 2011 from Equation (30) and the calibrated friction elasticity $\alpha^T \theta^T = 2$, we reconstruct

$$\tilde{\lambda}_{m|n}^T = \frac{\left(\widehat{b_{mt}^{\theta^T} / \bar{b}_t^{\theta^T}} \right) t_{nm}^{-\alpha^T \theta^T}}{\sum_{m' \in \mathcal{N} \setminus \{n\}} \left(\widehat{b_{m't}^{\theta^T} / \bar{b}_t^{\theta^T}} \right) t_{nm'}^{-\alpha^T \theta^T}}, \quad (\text{C.4})$$

where t_{nm} is the observed minimum bilateral travel time in 2011. We then use these fitted shares $\tilde{\lambda}_{m|n}^T$ in place of $\lambda_{m|n}^T$ when evaluating Equation (C.3). This reconstruction imposes the model’s own gravity structure on the baseline portfolio of leisure destinations, thereby reducing the influence of measurement error in the UnionPay microdata while preserving the estimated destination attractiveness and observed travel times.

Next, we estimate $\hat{b}_m^{\theta^T}$. However, in practice, leisure travel gravity estimation identifies destination attractiveness and origin accessibility only up to a common scale in each year. Centering the estimated destination–year fixed effects as in Equation (30), we recover the relative attractiveness measures $\widehat{b_{mt}^{\theta^T} / \bar{b}_t^{\theta^T}}$. Their intertemporal ratios therefore identify relative changes only up to the national geometric mean:

$$\frac{\exp \left(\tilde{\delta}_{m,2015} - \frac{1}{|\tilde{\mathcal{N}}_{2015}^{\text{des}}|} \sum_{m' \in \tilde{\mathcal{N}}_{2015}^{\text{des}}} \tilde{\delta}_{m',2015} \right)}{\exp \left(\tilde{\delta}_{m,2011} - \frac{1}{|\tilde{\mathcal{N}}_{2011}^{\text{des}}|} \sum_{m' \in \tilde{\mathcal{N}}_{2011}^{\text{des}}} \tilde{\delta}_{m',2011} \right)} \rightarrow_p \frac{\hat{b}_m^{\theta^T}}{\hat{\bar{b}}^{\theta^T}}.$$

Substituting these recovered $\widehat{b_{mt}^{\theta^T} / \hat{\bar{b}}^{\theta^T}}$ into Equation (C.3) yields the estimated accessibility change

(adjusted by a common scale $\hat{b}^{\theta^T}$):

$$\overline{\hat{\Phi}_n^T / \hat{b}^{\theta^T}} = \sum_{m \in \mathcal{N} \setminus \{n\}} \left(\overline{\hat{b}_m^{\theta^T} / \hat{b}^{\theta^T}} \right) \hat{t}_{nm}^{-\alpha^T \theta^T} \tilde{\lambda}_{m|n}^T.$$

Because $\ln \hat{b}^{\theta^T}$ does not vary across j and n , it is absorbed by the origin fixed effects in Equation (C.2). Consequently, using $\ln \overline{\hat{\Phi}_n^T / \hat{b}^{\theta^T}}$ as the regressor is equivalent to using $\Delta \ln \Phi_n^T$ up to an additive constant.

C.2.2 Instrumental variable

Even after first-differencing, $\Delta \ln \Phi_n^T$ may remain endogenous. In particular, cities that experienced larger HSR-induced travel time reductions over 2010–2015 may also have experienced correlated changes in local amenities or residual migration costs. Building on the same logic as in the main text, we construct an instrument that holds bilateral travel times fixed and retains only the contribution of changes in external destination attractiveness. Using the same reconstructed baseline shares, define

$$\tilde{\varphi}_n^T = \sum_{m \in \mathcal{N} \setminus \{n\}} \left(\overline{\hat{b}_m^{\theta^T} / \hat{b}^{\theta^T}} \right) \tilde{\lambda}_{m|n}^T. \quad (\text{C.5})$$

Relative to Equation (C.3), the instrument shuts down the $\hat{t}_{nm}^{-\alpha^T \theta^T}$ channel and therefore captures whether the portfolio of external destinations of city n became more attractive. The identifying assumption is that changes in other cities' leisure attractiveness are orthogonal to changes in city n 's unobserved residential amenities and residual migration costs, conditional on the origin fixed effects. We then instrument $\ln \overline{\hat{\Phi}_n^T / \hat{b}^{\theta^T}}$ with $\ln \tilde{\varphi}_n^T$ in the first-difference regression (C.2).

C.2.3 Estimation results

Table C.2 reports the IV estimates of Equation (C.2). Column (1) is the baseline first-difference specification. The coefficient on leisure travel accessibility is 0.144 and statistically significant at the 10% level. This point estimate is quantitatively close to the pooled IV estimate of 0.180 in Column (6) of Table 4, which we adopt as the baseline value of $\theta^R \psi / \theta^T$ in the main text. The similarity of the two estimates suggests that the migration response to leisure accessibility is not an artifact of cross-sectional differences in accessibility levels, and therefore provides a reassuring robustness check for the main-text parameter.

At the same time, we find the first-difference estimate is less precisely estimated. Relative to the pooled levels regression, the sample shrinks substantially (from 28,446 to 7,806 observations): any city pair with a zero migration share in either census year drops out of the log-differenced sample. Moreover, first-differencing removes the large cross-sectional component of accessibility variation and retains only within-city changes between 2010 and 2015, which are quantitatively smaller. Both forces inflate the standard error, even though the point estimate itself remains comparable to the main-text benchmark.

A natural concern with interpreting the accessibility coefficient as a leisure-amenity channel is that HSR may also lower bilateral migration costs directly by compressing passenger travel time between origin and destination. Although a one-time moving fare is unlikely to matter for long-run location choice, recurrent home visits can directly enter residential utility as a term that is discounted by bilateral travel time, rather than entering the leisure travel decision. In reduced form, this term is isomorphic to a component of migration costs. To assess this alternative mechanism, Column (2) augments the first-difference regression with the change in bilateral travel time, $\Delta \ln t_{jn}$. If this is an important omitted channel, one may expect $\Delta \ln t_{jn}$ to enter with a meaningfully negative coefficient and to absorb part of the accessibility effect. Empirically, the travel-time coefficient is small and statistically insignificant, while the accessibility estimate remains unchanged. Therefore, it provides little evidence that the home-tie mechanism drives the estimated leisure-accessibility coefficient.

Table C.2: Estimation of the first-difference migration gravity equation

	Dependent variable: $\Delta \ln \lambda_{jn} - \theta^R \Delta \ln \frac{w_n}{P_n}$	
	(1)	(2)
$\ln \widehat{\phi_n^T} / \widehat{b}^{\theta^T}$	0.144* (0.079)	0.144* (0.078)
$\Delta \ln t_{jn}$		-0.014 (0.078)
Origin fixed effects	Yes	Yes
Estimator	IV	IV
First-stage F-statistic	9,885	9,800
# of observations	7,806	7,806

Notes: This table reports IV estimates of the first-difference migration gravity equation. Column (1) estimates Equation (C.2). Column (2) additionally controls for the change in bilateral travel time. Kleibergen–Paap first-stage F-statistics are reported. Robust standard errors are reported in parentheses.

*** p < 0.01, ** p < 0.05, * p < 0.1.

C.3 Goodness of fit of the leisure travel and migration gravity equations

This appendix evaluates how well the estimated leisure travel and migration gravity equations reproduce two aggregate moments that are not directly targeted in estimation: city-level inbound leisure travel volumes and city-level residential populations.

C.3.1 Leisure travel gravity

Recall that the leisure travel gravity equation identifies the composite friction elasticity $\alpha^T \theta^T$ together with origin and destination fixed effects. Given the recovered relative destination attractiveness $\widehat{b_{mt}^{\theta^T}} / \widehat{b_t^{\theta^T}}$ from Equation (30) and the calibrated friction elasticity $\alpha^T \theta^T = 2$, the model-implied leisure travel shares are

$$\tilde{\lambda}_{m|n,t}^T = \frac{\left(\widehat{b_{mt}^{\theta^T}} / \widehat{b_t^{\theta^T}} \right) t_{nmt}^{-\alpha^T \theta^T}}{\sum_{m' \in \mathcal{N} \setminus \{n\}} \left(\widehat{b_{m't}^{\theta^T}} / \widehat{b_t^{\theta^T}} \right) t_{nm't}^{-\alpha^T \theta^T}}, \quad m \in \mathcal{N} \setminus \{n\}. \quad (\text{C.6})$$

The common normalizing factor $\bar{b}_t^{\theta^T}$ cancels in the ratio, thus this prediction is completely consistent with the structural equation of $\lambda_{m|n,t}^T$ in Equation (3). To assess fit at the city level, we take each city's observed total outbound leisure travel volume, $Volume_{nt}^{\text{out}}$, as given, and predict its total inbound volume by reallocating these outbound trips across destinations according to Equation (C.6):

$$\widehat{Volume}_{mt}^{\text{in}} = \sum_{n \in \mathcal{N} \setminus \{m\}} \tilde{\lambda}_{m|n,t}^T Volume_{nt}^{\text{out}}. \quad (\text{C.7})$$

Therefore, Equation (C.7) asks whether travel times and the estimated destination attractiveness are sufficient to redistribute the observed origin totals into a destination pattern that matches the data.

Figure C.1 compares the predicted inbound volume with the observed inbound volume in 2011. The predicted and observed volumes are closely aligned, indicating that the leisure travel gravity equation accounts for the cross-sectional distribution of inbound leisure flows well, even though the estimating equation includes only bilateral travel time and two-way fixed effects. Residual discrepancies are modest: less attractive destinations tend to be slightly under-predicted, whereas more attractive destinations tend to be slightly over-predicted.




Figure C.1: Predicted and observed inbound leisure travel volumes in 2011

Notes: This figure compares the predicted and observed inbound leisure travel volumes in 2011. Each point represents a destination city. Observed inbound volumes are obtained by aggregating UnionPay leisure trips across origins. Predicted inbound volumes are constructed from Equation (C.7), which reallocates each origin's observed outbound volume according to the model-implied leisure travel shares in Equation (C.6). The dashed line is the 45-degree line. The reported uncentered R-squared is obtained from a regression of the predicted volume on the observed volume without a constant.

C.3.2 Migration gravity

Different from the estimated leisure travel gravity equation, our estimation of the migration gravity does not recover the full structural specification of the migration probabilities. As discussed in the main text, we estimate

$$\ln \lambda_{jnt} - \theta^R \ln \frac{w_{nt}}{P_{nt}} = \theta^R \frac{\psi}{\theta^T} \ln \frac{\Phi_{nt}^T}{\bar{\Phi}_t^T} + \beta_t \ln (\text{Distance}_{jn} + 1) + \delta_{jt} + \epsilon_{jnt},$$

where the calibrated real-wage term is moved to the left-hand side, bilateral distance enters as a proxy for distance-related migration costs, and origin–year fixed effects absorb multilateral resistance. The composite residual ϵ_{jnt} collects unobserved residential amenities and residual components of migration frictions. Consequently, simply predicting the left-hand side of this regression and adding back $\theta^R \ln(w_{nt}/P_{nt})$ does not yield a well-defined probability system: the resulting objects need not lie in $(0, 1)$ and need not sum to one across destinations for each origin j .

We therefore construct predicted migration shares from the structural definition in Equation (5), retaining only the components that are either observed or identified in estimation. Specifically, we use observed real wages, the recovered relative leisure accessibility $\widehat{\Phi_{nt}^T}/\bar{\Phi}_t^T$, and the estimated distance elasticity, while shutting down unobserved amenities and residual bilateral migration costs:

$$\tilde{\lambda}_{jnt} = \frac{w_{nt}^{\theta^R} P_{nt}^{-\theta^R} \left(\widehat{\Phi_{nt}^T}/\bar{\Phi}_t^T \right)^{\theta^R \psi / \theta^T} (\text{Distance}_{jn} + 1)^{\beta_t}}{\sum_{n' \in \mathcal{N}} w_{n't}^{\theta^R} P_{n't}^{-\theta^R} \left(\widehat{\Phi_{n't}^T}/\bar{\Phi}_t^T \right)^{\theta^R \psi / \theta^T} (\text{Distance}_{jn'} + 1)^{\beta_t}}. \quad (\text{C.8})$$

Compared to Equation (5), this construction imposes $a_{nt} = 1$ for all n and replaces $\iota_{jnt}^{-\theta^R}$ with the distance term $(\text{Distance}_{jn} + 1)^{\beta_t}$. The common normalizing factor $\bar{\Phi}_t^T$ cancels in the ratio and therefore does not affect $\tilde{\lambda}_{jnt}$. By construction, $\tilde{\lambda}_{jnt} \in (0, 1)$ and $\sum_{n \in \mathcal{N}} \tilde{\lambda}_{jnt} = 1$ for every origin j . Given the initial *hukou* labor endowment $\{\bar{L}_j\}_{j \in \mathcal{N}}$, the predicted number of residents in destination n is then

$$\tilde{L}_{nt} = \sum_{j \in \mathcal{N}} \tilde{\lambda}_{jnt} \bar{L}_j. \quad (\text{C.9})$$

Figure C.2 compares the predicted residential population with the observed one in 2010. Because city-level wage data are missing for a subset of cities, the comparison is restricted to the city pairs for which all objects entering Equation (C.8) are available. As shown in the figure, predicted and observed populations remain highly correlated, indicating that real wages, leisure accessibility, and distance jointly capture an important component of the spatial distribution of residents.

C.4 Recovering production-labor shares from the MRIO table

Constructing the denominator, $\sum_{k=1}^K \pi_{jn}^k D_n^k$, in the business travel volume regression requires the sectoral production-labor shares, defined as $v_n^s \equiv L_n^s/L_n$. The MRIO tables for 2012 and 2015 report the sectoral compensation of labor. While these wage bills identify the allocation of production labor across sectors within a given city, they do not capture the allocation relative to total city employment,




Figure C.2: Predicted and observed residential populations in 2010

Notes: This figure compares the predicted and observed residential populations in 2010. Each point represents a destination city. Observed populations are obtained by aggregating census migration flows across origins. Predicted populations are constructed from Equation (C.9), using the incomplete structural migration shares in Equation (C.8), which retain observed real wages, recovered leisure accessibility, and the estimated distance elasticity, while shutting down unobserved amenities and residual non-distance migration costs. Because wage data are missing for some cities, the sample is restricted to city pairs for which all objects entering Equation (C.8) are available. The dashed line is the 45-degree line. The reported uncentered R-squared is obtained from a regression of the predicted population on the observed population without a constant.

L_n , because the latter also encompasses labor hired by final goods firms for search and entry activities. Accordingly, we define $\chi_n^k \equiv L_n^k / \sum_{s=1}^K L_n^s$ as sector k 's share of city n 's aggregate production labor. Because our structural model assumes perfect labor mobility across sectors within a city, thereby equalizing wages across all local sectors, the observed share of labor compensation for sector k in the MRIO data maps identically to its employment share, χ_n^k .

To map χ_n^k into ν_n^k , we combine the goods and labor market clearing conditions. Recall that total final-goods expenditure in sector k of city n is

$$X_n^k = \rho_n^k w_n L_n + \sum_{s=1}^K \gamma_n^{s,k} R_n^s,$$

and that labor market clearing requires

$$w_n L_n = \sum_{k=1}^K \left(\gamma_n^k R_n^k + \frac{1}{\sigma} X_n^k \right).$$

Production labor satisfies $w_n L_n^k = \gamma_n^k R_n^k$. Moreover, because intermediate producers in sector s allocate cost shares that sum to one, $\gamma_n^s + \sum_{k=1}^K \gamma_n^{s,k} = 1$, and therefore $\sum_{k=1}^K \gamma_n^{s,k} = 1 - \gamma_n^s$. Aggregating final-goods expenditure across sectors then yields

$$\begin{aligned} \sum_{k=1}^K X_n^k &= w_n L_n + \sum_{s=1}^K (1 - \gamma_n^s) R_n^s \\ &= w_n L_n + \sum_{s=1}^K \frac{w_n L_n^s}{\gamma_n^s} - \sum_{s=1}^K w_n L_n^s. \end{aligned}$$

Substituting into the labor market clearing condition gives

$$w_n L_n = \sum_{k=1}^K w_n L_n^k + \frac{1}{\sigma} \left(w_n L_n + \sum_{s=1}^K \frac{w_n L_n^s}{\gamma_n^s} - \sum_{s=1}^K w_n L_n^s \right).$$

Dividing through by w_n and rearranging, we obtain the exact relationship between total employment and production labor:

$$L_n = \sum_{k=1}^K L_n^k + \frac{1}{\sigma - 1} \sum_{k=1}^K \frac{L_n^k}{\gamma_n^k}. \quad (\text{C.10})$$

The first term on the right-hand side is aggregate production labor; the second term is the non-production labor hired for search and entry, expressed as a function of the production-labor vector and the labor production elasticity. Substituting $L_n^k = \chi_n^k \sum_{s=1}^K L_n^s$ into Equation (C.10) yields

$$L_n = \left(1 + \frac{1}{\sigma - 1} \sum_{k=1}^K \frac{\chi_n^k}{\gamma_n^k} \right) \sum_{s=1}^K L_n^s,$$

and therefore

$$\sum_{s=1}^K L_n^s = \left(1 + \frac{1}{\sigma - 1} \sum_{k=1}^K \frac{\chi_n^k}{\gamma_n^k} \right)^{-1} L_n.$$

The production-labor share relative to total employment is then

$$\nu_n^s = \frac{L_n^s}{L_n} = \left(1 + \frac{1}{\sigma - 1} \sum_{k=1}^K \frac{\chi_n^k}{\gamma_n^k} \right)^{-1} \chi_n^s. \quad (\text{C.11})$$

We recover $\{\nu_n^k\}_{n \in \mathcal{N}, k \in \mathcal{K}}$ based on Equation (C.11) using the calibrated elasticity of substitution $\sigma = 6$ together with the MRIO-based $\{\chi_n^k, \gamma_n^k\}_{n \in \mathcal{N}, k \in \mathcal{K}}$, and then construct

$$D_n^k = \rho_n^k + \sum_{s=1}^K \frac{\gamma_n^{s,k}}{\gamma_n^s} \nu_n^s$$

for the business travel volume regression. We separately use the 2012 and 2015 city-level MRIO tables to construct D_n^k for the two years. Specifically, for 2015, not only $\{\chi_n^s\}_{n \in \mathcal{N}, s \in \mathcal{S}}$ but also $\{\rho_n^k, \gamma_n^s, \gamma_n^{s,k}\}_{n \in \mathcal{N}, s, k \in \mathcal{K}}$ are reconstructed based on the 2015 MRIO table.

D Appendix for counterfactual analysis

D.1 Travel time decomposition

The decomposition in Equation (40) requires the relative change in each pair's volume share, $\hat{s}_{nj}$. We construct this object separately for business and leisure travel.

D.1.1 Business travel

Total business travel volume mixes procurement-related travel with residual trips. Following Section 4.2.3, we write

$$V_{nj}^{\text{business}} = (1 + \phi_{nj}) \sum_{k=1}^K F_n^k l_{nj}^{k,f},$$

where the second term is the search labor implied by Equation (34) and ϕ_{nj} is the ratio of residual to procurement-related travel. The relative change in observed volume is therefore the product of the two components,

$$\hat{V}_{nj}^{\text{business}} = \widehat{(1 + \phi_{nj})} \times \sum_{k=1}^K \widehat{F_n^k l_{nj}^{k,f}}.$$

The counterfactual is designed to isolate the procurement margin. The structural model determines search labor, trade shares, and the distribution of labor, but not ϕ_{nj} . We therefore set $\widehat{1 + \phi_{nj}} = 1$ and compute

$$\hat{V}_{nj}^{\text{business}} = \sum_{k=1}^K \widehat{F_n^k l_{nj}^{k,f}}.$$

This is not a claim that residual travel is invariant to HSR; instead, the resulting decomposition answers how national average business travel time changes when procurement-related volumes adjust, holding the residual-to-procurement ratio fixed. Shocks that operate through ϕ_{nj} are shut down by construction in our counterfactual simulation and residual travel scales with procurement-related travel.

With ϕ_{nj} fixed, the relative change in the volume share is then

$$\hat{s}_{nj}^{\text{business}} \equiv \frac{\hat{V}_{nj}^{\text{business}}}{\sum_{n'=1}^N \sum_{j'=1}^N \hat{V}_{n'j'}^{\text{business}} s_{n'j'}^{\text{business}}} = \frac{\sum_{k=1}^K \widehat{F_n^k l_{nj}^{k,f}}}{\sum_{n'=1}^N \sum_{j'=1}^N s_{n'j'}^{\text{business}} \sum_{k=1}^K \widehat{F_{n'}^k l_{n'j'}^{k,f}}}.$$

Substituting Equation (34), we further obtain

$$\hat{s}_{nj}^{\text{business}} = \frac{\hat{t}_{nj}^{-\alpha^P} \sum_{k=1}^K \hat{\pi}_{jn}^k \hat{L}_n^k \varpi_{jn}^k}{\sum_{n'=1}^N \sum_{j'=1}^N s_{n'j'}^{\text{business}} \hat{t}_{n'j'}^{-\alpha^P} \sum_{k=1}^K \hat{\pi}_{j'n'}^k \hat{L}_{n'}^k \varpi_{j'n'}^k},$$

where $\varpi_{jn}^k \equiv \pi_{jn}^k \tilde{L}_n^k / \sum_{k'=1}^K \pi_{jn}^{k'} \tilde{L}_n^{k'}$ is calculated based on variables in the calibrated no-HSR baseline equilibrium, and all relative changes are calculated based on the solved counterfactual equilibrium. With the observed shares $\{s_{nj}^{\text{business}}, S_{nj}^{\text{business}}\}_{n,j \in \mathcal{N}}$ from 2011 UnionPay data, we calculate $\hat{s}_{nj}^{\text{business}}$

for each city pair. This further allows us to calculate all three components in Equation (40).

The corresponding city-level average business travel time is $\hat{t}_n^{\text{business}} = \sum_{j=1}^N t_{nj} s_{j|n}^{\text{business}}$, with $s_{j|n}^{\text{business}} \equiv s_{nj}^{\text{business}} / \sum_{j'=1}^N s_{nj'}^{\text{business}}$. Its relative change is

$$\hat{t}_n^{\text{business}} = \sum_{j=1}^N s_{j|n}^{\text{business}} \hat{t}_{nj} \hat{s}_{j|n}^{\text{business}}, \quad (\text{D.1})$$

where $s_{j|n}^{\text{business}} \equiv s_{nj}^{\text{business}} / \sum_{j'=1}^N s_{nj'}^{\text{business}}$ and

$$\hat{s}_{j|n}^{\text{business}} = \frac{\hat{t}_{nj}^{-\alpha^P} \sum_{k=1}^K \hat{\pi}_{jn}^k \hat{L}_n^k \varpi_{jn}^k}{\sum_{j'=1}^N s_{j'|n}^{\text{business}} \hat{t}_{nj'}^{-\alpha^P} \sum_{k=1}^K \hat{\pi}_{j'n}^k \hat{L}_n^k \varpi_{j'n}^k}.$$

D.1.2 Leisure travel

Leisure volume from n to m can be written as $V_{nm}^{\text{leisure}} = \lambda_{m|n}^T \text{OutTraveler}_n$, where OutTraveler_n is the number of intercity leisure travelers originating in n . The model determines $\hat{\lambda}_{m|n}^T$ but not $\widehat{\text{OutTraveler}}_n$, because the extensive margin of outbound leisure travel is not specified. The national share change

$$s_{nm}^{\text{leisure}} = \frac{\hat{\lambda}_{m|n}^T \widehat{\text{OutTraveler}}_n}{\sum_{n'=1}^N \sum_{m' \neq n'} s_{n'm'}^{\text{leisure}} \hat{\lambda}_{m'|n'}^T \widehat{\text{OutTraveler}}_{n'}}$$

is therefore not identified, and neither are the national composition, interaction, or total effects for leisure travel. Direct time effects remain identified, since they depend only on observed s_{nm}^{leisure} and $\hat{t}_{nm}$.

Although the national composition and interaction terms are not identified, the change in average outbound leisure time from a given city, $\hat{t}_n^{\text{leisure}}$, is identified. Conditional on origin n , we have

$$\hat{s}_{m|n}^{\text{leisure}} \equiv \frac{\hat{V}_{nm}^{\text{leisure}}}{\sum_{m' \neq n} s_{m'|n}^{\text{leisure}} \hat{V}_{nm'}^{\text{leisure}}} = \frac{\hat{\lambda}_{m|n}^T}{\sum_{m' \neq n} s_{m'|n}^{\text{leisure}} \hat{\lambda}_{m'|n}^T},$$

where $s_{m|n}^{\text{leisure}} \equiv V_{nm}^{\text{leisure}} / \sum_{m' \neq n} V_{nm'}^{\text{leisure}}$. Therefore,

$$\hat{t}_n^{\text{leisure}} = \sum_{m \neq n} s_{m|n}^{\text{leisure}} \hat{t}_{nm} \hat{s}_{m|n}^{\text{leisure}}, \quad (\text{D.2})$$

where $s_{m|n}^{\text{leisure}} \equiv s_{nm}^{\text{leisure}} / \sum_{m' \neq n} s_{nm'}^{\text{leisure}}$ is constructed based on observed shares.

D.1.3 Comparison with observed business travel, 2011–2015

The counterfactuals in the main text are defined relative to the no-HSR baseline. To compare the model with observed travel patterns, we instead express the 2015 equilibrium relative to 2011, the first year of the UnionPay sample. Specifically, we solve the counterfactual for each year and normalize, so that, for example, $\hat{\pi}_{jn,2015/2011}^k = \hat{\pi}_{jn,2015}^k / \hat{\pi}_{jn,2011}^k$. Observed time and volume changes are computed from the same 2011–2015 window.

Table D.1 reports the resulting decomposition of national average business travel time. The direct time effects are close (-16.816% in the data versus -13.889% in the model): HSR accounts for most of the observed time effects over this period. The composition effects are not. The data exhibit a large shift toward longer routes (22.135%), whereas the model, which isolates the HSR shock, generates only a modest reallocation (4.888%). This discrepancy is expected. Observed travel shares respond to many concurrent forces—regional policy, industrial relocation, and income growth—that lie outside the counterfactual. By contrast, the interaction effects match. The model therefore reproduces the covariance between travel-time reductions and the reallocation of business travel.

Table D.1: Percentage change in national average business travel time, 2011–2015

	Data	Model
Total effects	3.234	−13.020
Direct time effects	−16.816	−13.889
Composition effects	22.135	4.888
Interaction effects	−4.091	−4.018

Notes: This table decomposes the change in national average business travel time between 2011 and 2015 using Equation (40), in percent. The model column reports the 2015 counterfactual expressed relative to 2011, as described in the text. Baseline shares are measured from 2011 UnionPay business volumes. In the data column, the three components need not sum to the total effect. Equation (40) is weighted by 2011 time shares S_{nj} and requires $\hat{s}_{nj} = s'_{nj}/s_{nj}$ to be defined. City pairs with zero UnionPay volume in 2011 therefore drop out of the direct time, composition, and interaction terms, but they still enter the 2015 national average that defines the total effect. The residual is this extensive-margin contribution. The model column constructs $\hat{s}_{nj}$ from the structural formula without dividing by baseline shares, so the identity holds.

D.2 The role of endogenous search and other equilibrium responses

The channel decompositions in the main text vary whether HSR-induced travel-time reductions enter the business-search or leisure-travel block, while allowing all endogenous variables to adjust in equilibrium. We now conduct a complementary set of exercises that restrict particular equilibrium responses conditional on these shocks. We first construct a coherent no-search benchmark to quantify the incremental role of endogenous supplier search. We then report two diagnostic exercises that hold migration shares or sourcing shares fixed. However, we highlight that these exercises are not an additive decomposition of the full effect because prices, wages, and other endogenous variables are re-solved under each restricted system.

D.2.1 A no-search Eaton–Kortum benchmark

We first replace the endogenous-search block with a standard Eaton–Kortum sourcing block while retaining the remaining structure of the model, including input–output linkages, final goods firm entry, housing, leisure accessibility, and migration. In the absence of search frictions, final goods

firms observe the complete set of potential suppliers. The trade share is therefore

$$\pi_{jn}^{k,NS} = \frac{T_j^k \left(c_j^{k,NS} \bar{d}_{jn} \right)^{-\theta^P}}{\sum_{j'=1}^N T_{j'}^k \left(c_{j'}^{k,NS} \bar{d}_{j'n} \right)^{-\theta^P}},$$

where the superscript NS denotes the no-search benchmark. Taking the observed baseline trade shares as sufficient statistics, the relative change satisfies

$$\hat{\pi}_{jn}^{k,NS} = \frac{\left(\hat{c}_j^{k,NS} \right)^{-\theta^P}}{\sum_{j'=1}^N \left(\hat{c}_{j'}^{k,NS} \right)^{-\theta^P} \pi_{j'n}^k}. \quad (D.3)$$

Unlike Equation (A.5), this expression contains neither the endogenous-search amplification of the trade elasticity nor the passenger travel-time term.

Furthermore, the unit cost of the final good is

$$c_n^{k,NS} = \gamma \left[\sum_{j=1}^N T_j^k \left(c_j^{k,NS} \bar{d}_{jn} \right)^{-\theta^P} \right]^{-\frac{1}{\theta^P}},$$

and its relative change is

$$\hat{c}_n^{k,NS} = \left[\sum_{j=1}^N \left(\hat{c}_j^{k,NS} \right)^{-\theta^P} \pi_{jn}^k \right]^{-\frac{1}{\theta^P}}.$$

Because firms incur no search expenditure in this benchmark, gross profits are exhausted solely by the fixed entry cost. Free entry implies

$$F_n^{k,NS} = \frac{1}{\sigma} \frac{\tilde{L}_n^{k,NS}}{\bar{c}_n^k},$$

which differs in level from the corresponding expression in the endogenous-search model. However, the coefficients and fixed entry costs are invariant across equilibria within each model, so both models imply

$$\hat{F}_n^k = \hat{L}_n^k.$$

Given c_n^k , the removal of endogenous search does not affect the optimal price of monopolistic final goods firms and the aggregate rule for final-good prices. Therefore, the sector-specific final-good price index in the no-search benchmark still satisfies $p_n^{k,NS} = \sigma/(\sigma-1) c_n^{k,NS} (F_n^{k,NS})^{1/(1-\sigma)}$. Therefore, the relative change is given by

$$\hat{p}_n^{k,NS} = \left(\hat{L}_n^{k,NS} \right)^{\frac{1}{1-\sigma}} \left[\sum_{j=1}^N \left(\hat{c}_j^{k,NS} \right)^{-\theta^P} \pi_{jn}^k \right]^{-\frac{1}{\theta^P}}. \quad (D.4)$$

The labor-market clearing equation also remains unchanged:

$$w_n^{\text{NS}} L_n^{\text{NS}} = \sum_{k=1}^K \gamma_n^k R_n^{k,\text{NS}} + \frac{1}{\sigma} \sum_{k=1}^K X_n^{k,\text{NS}}.$$

In the full model, the second term combines labor payments for search and entry; in the no-search benchmark, it consists entirely of entry labor.

It is straightforward to verify that other relative-change equations defined in Definition A.1 are unchanged. Therefore, Equations (D.3) and (D.4) together with all other unchanged relative-change equations constitute a no-search spatial equilibrium in relative changes. We use the same algorithm to loop through the nonlinear system until it converges. In this case, passenger travel time continues to affect household leisure accessibility but no longer enters firms' sourcing decisions.

D.2.2 Fixed migration shares

To examine the role of labor reallocation, we hold every bilateral migration share at its baseline value in the *endogenous-search full model*. In the algorithm, we set $\hat{\lambda}_{jn} = 1$ for all $j, n \in \mathcal{N}$ and loop until the system converges. Therefore, all business-search, trade, price, and market-clearing equations remain active. The resulting real-wage effect measures the production-side consequences of HSR when the spatial distribution of labor is held fixed.

Equation (42) cannot be applied to this exercise because it relies on workers optimally reselecting their residential locations. Therefore, we only use the fixed-migration exercise to interpret real-wage effects and do not include its welfare index in the welfare comparisons below.

D.2.3 Fixed trade shares

As a further diagnostic, we hold bilateral sectoral trade shares fixed in the *no-search benchmark model*. Specifically, we set $\hat{\pi}_{jn}^k = 1$ for all $j, n \in \mathcal{N}$ and $k \in \mathcal{K}$ in the loop. This exercise prevents final goods firms from both searching and reallocating expenditure across supplier cities. Production costs, wages, entry, output, and migration nevertheless continue to adjust through the baseline trade network.

Although the trade network adjustment is manually muted, the import prices of intermediate goods still change with the exporter's marginal cost. For a variety ω^k sourced from city j in the baseline, its change in import price in city n is given by

$$\hat{c}_{jn}^{k,\text{FT}}(\omega^k) = \hat{c}_j^{k,\text{FT}},$$

where the superscript FT denotes fixed trade. We approximate the change in the importing firm's unit cost in city n using the baseline trade shares:

$$\hat{c}_n^{k,\text{FT}} \approx \sum_{j=1}^N \pi_{jn}^k \hat{c}_j^{k,\text{FT}}.$$

The corresponding change in the final-good price index in city n is then

$$\hat{p}_n^{k,FT} = \left(\hat{L}_n^{k,FT} \right)^{\frac{1}{1-\sigma}} \hat{c}_n^{k,FT}. \quad (\text{D.5})$$

Finally, we replace Equation (D.4) in Section D.2.1 with Equation (D.5) and loop until the system converges, from which we approximate the results when both search and trade responses are muted.

D.2.4 Results

Table D.2 reports the 2015 results. The full equilibrium allows both residential and production-network responses, which is consistent with the counterfactual in the main text. The following three columns report the results of the above restricted exercises.

Table D.2: The role of endogenous search, migration, and trade responses

	Full equilibrium	Fixed migration	No-search benchmark	Fixed trade
<i>Panel A. Change in national average real wage</i>				
Both channels	2.185	0.800	0.710	0.356
Business-search channel	1.092	0.800	0.000	0.000
Leisure channel	0.951	0.000	0.710	0.356
<i>Panel B. Change in national welfare</i>				
Both channels	3.140		2.658	2.773
Business-search channel	0.554		0.000	0.000
Leisure channel	2.574		2.658	2.773

Notes: This table reports counterfactual outcomes under the 2015 HSR network relative to the no-HSR baseline. Panel A reports the percentage change in the national average real wage, and Panel B reports the log change in national welfare. “Both channels” allows HSR-induced travel-time reductions to affect both leisure travel and business search. The search-only and leisure-only counterfactuals hold the other friction at its baseline level. The “Full equilibrium” column replicates the results in the main text. The fixed-migration exercise holds the spatial distribution of labor constant. Since the fixed-migration exercise does not optimize residential choice, the relevant welfare index differs from the welfare measure in the main text and is not reported. The no-search benchmark retains endogenous trade-share adjustment through a standard Eaton–Kortum sourcing block, as defined in Section D.2.1, whereas the fixed-trade exercise further holds trade shares at their baseline values and uses Equation (D.5) for the final-good price change.

Among these exercises, the comparison between the full equilibrium and the no-search benchmark speaks most directly to the contribution of endogenous supplier search. In the full model, the 2015 HSR network raises the national average real wage by 2.185%. Replacing the endogenous-search block with a standard Eaton–Kortum sourcing block, while retaining input–output linkages, firm entry, housing, leisure accessibility, and migration, reduces this gain to 0.71%. In the conventional benchmark, passenger travel time has no direct effect on firms’ sourcing decisions because HSR does not alter freight costs. The remaining gain is generated by the leisure-accessibility shock, the induced labor reallocation, and their general-equilibrium effects within the standard production structure. Endogenous search therefore raises the predicted real-wage effect by more than a factor of three. However, the difference between the two predictions incorporates both the direct business-search effect and its interactions with other equilibrium responses, and should therefore be

interpreted as a comparison between models rather than as an additive contribution of search.

The fixed-migration exercise provides a complementary diagnostic within the full endogenous-search model. When migration shares are held constant, the real-wage effect of the business-search channel declines from 1.092% to 0.8%. Thus, most of the production-side gain survives without labor reallocation, while migration feedback further amplifies it. By contrast, the leisure channel produces no real-wage gain when migration is fixed. This result follows directly from the model structure: leisure accessibility does not enter firms' production decisions and can affect real income only by changing the spatial allocation of labor.

The leisure channel also illustrates how a consumption-amenity shock propagates through the production system once migration is allowed. Its real-wage effect is 0.951% in the full model, 0.71% in the no-search benchmark, and 0.356% when bilateral trade shares are additionally held fixed. The final exercise allows wages, marginal costs, firm entry, and output to adjust through the baseline trade network. Allowing trade shares to respond according to the standard Eaton–Kortum sourcing rule approximately doubles the real-wage gain from 0.356% to 0.71%, while endogenous search raises it further to 0.951%. Hence, endogenous search amplifies not only the direct business-search shock but also the production consequences of leisure-induced labor reallocation.

Welfare is considerably less sensitive to the presence of endogenous search. The no-search benchmark retains approximately 85% of the full welfare gain. The fixed-trade diagnostic similarly produces a welfare gain of 2.773%. These results are consistent with the decomposition in Equation (44): improved leisure accessibility enters household utility directly and therefore continues to generate substantial welfare gains even when the business-search mechanism and endogenous trade-share adjustment are absent.

Taken together, the no-search benchmark shows that endogenous supplier search is central to the real-income effects of HSR. The restricted-response exercises further show that migration amplifies the business-search effect, while trade-share adjustment and endogenous search amplify the production consequences of leisure-induced labor reallocation.

D.3 Spatial inequality in residence-based welfare

The national welfare measure in the main text aggregates origin-specific expected utility and is designed to evaluate changes in welfare rather than its cross-sectional distribution in levels. Constructing a welfare Gini requires baseline welfare levels across cities, which the exact-hat counterfactual system does not recover. In particular, we do not observe the baseline levels of local amenities. We therefore conduct a supplementary calculation of the spatial dispersion of a residence-based welfare index under additional assumptions.

Conditional on residing in city n , the deterministic component of utility is proportional to

$$\mathcal{W}_n = a_n \frac{w_n}{P_n} \left(\Phi_n^T \right)^{\psi/\theta^T},$$

where a_n denotes local amenities, w_n/P_n is the local real wage, and Φ_n^T measures nonlocal leisure accessibility. This index describes the utility associated with residing in a location before the realiza-

tion of idiosyncratic residential preferences. It is therefore distinct from the origin-specific expected utility in Equation (42), which incorporates the option to choose among residential destinations.

To construct its baseline level, we make three approximations. First, we assume that local amenities are identical across cities, $a_n = \bar{a}$. The common amenity term then has no effect on the Gini coefficient. Second, because baseline city-level consumer price levels are not recovered, we proxy P_n using the corresponding cumulative provincial consumer price index through 2011. Third, we recover relative leisure accessibility in 2011 from the origin fixed effects of the leisure-travel gravity equation and apply the inverse HSR shock to obtain its value in the no-HSR baseline. Denote the resulting relative accessibility measure by $\tilde{\Phi}_n^{T,0}$. The normalization of the fixed effect estimates is immaterial for the Gini coefficient.

Under these assumptions, the residence-based welfare index in the no-HSR baseline is

$$\mathcal{W}_n = \frac{w_n}{P_n} \left(\tilde{\Phi}_n^{T,0} \right)^{\psi/\theta^T}, \quad (\text{D.6})$$

In the counterfactual equilibrium, the corresponding index is

$$\mathcal{W}'_n = \frac{w_n \hat{w}_n}{P_n \hat{P}_n} \left(\tilde{\Phi}_n^{T,0} \hat{\Phi}_n^T \right)^{\psi/\theta^T}. \quad (\text{D.7})$$

Let $\text{Gini}(\cdot)$ denote the Gini coefficient across cities. The change in the spatial welfare Gini is given by

$$\widehat{\text{Gini}}^{\text{welfare}} - 1 = \frac{\text{Gini} \left(\{\mathcal{W}'_n\}_{n=1}^N \right)}{\text{Gini} \left(\{\mathcal{W}_n\}_{n=1}^N \right)} - 1, \quad (\text{D.8})$$

which measures the geographic dispersion of residence-based welfare opportunities across cities.

Table D.3 reports the comparison between the historical 2015 network and the universal network. The historical network raises the cross-city welfare Gini by 1.363%, while the universal network raises it by 1.236%. Universal connectivity therefore modestly attenuates the additional spatial welfare dispersion generated by HSR, although it does not eliminate it. In the single-channel counterfactuals, the contribution of leisure accessibility falls from 0.85% under the historical network to 0.568% under universal connectivity. The contribution of business search instead rises from 0.359% to 0.53%.

Table D.3: Spatial inequality in residence-based welfare under alternative HSR networks

	Network in 2015	Universal network
Both channels	1.363	1.236
Search channel	0.359	0.530
Leisure channel	0.850	0.568

Notes: This table reports the percentage change in the cross-city Gini coefficient of the residence-based welfare index defined in Equations (D.6) and (D.7). The results are measured relative to the no-HSR baseline and are not changes in Gini points. “Network in 2015” uses the historical HSR network observed in 2015, while “Universal network” connects every city pair using the travel-time construction described in the text. The search-only and leisure-only counterfactuals hold the other travel friction at its no-HSR level.

Table D.4 further shows the results of travel-prioritization network expansions. The change in the welfare Gini is non-monotonic in network coverage. It rises from 0.855% when the top 0.1% of links are connected to 1.472% at the top 1%, and reaches a maximum of 2.029% at the top 5%. It then declines as coverage expands, reaching 1.236% under universal connectivity. Early construction concentrates accessibility and real-income gains among a relatively small set of locations, while broader expansion spreads accessibility improvements across a larger set of cities and reduces their spatial dispersion. This decline occurs even as national welfare continues to increase.

Table D.4: Spatial inequality in residence-based welfare under alternative HSR networks

	Top percentage of links						
	0.1	1	5	10	20	50	100
Spatial welfare Gini	0.855	1.472	2.029	1.690	1.610	1.355	1.236

Notes: This table reports the percentage change in the cross-city Gini coefficient of the residence-based welfare index defined in Equations (D.6) and (D.7). It reports both-channel results for networks that connect city pairs in descending order of bilateral business-travel volume.

Again, we highlight that, since this regional-based measure does not account for the spatial reallocation of labor and omits cross-city differences in baseline local amenities, it should be interpreted as suggestive rather than as estimates of interpersonal welfare inequality.

E Additional figures and tables




Figure E.1: Evolution of the HSR network (Continued on the next page)





Figure E.1: Evolution of the HSR network (Continued)

Notes: The figure plots the HSR city-pair network for 2008–2015. A link is shown when a city pair is classified as connected by HSR in the processed HSR travel-time dataset. Directed origin–destination pairs are collapsed into undirected links, so each connected city pair appears at most once in each year. Lines are drawn as straight segments between population-weighted city centers and therefore represent intercity HSR connectivity rather than physical rail alignments.

(A) Relative leisure travel attractiveness in 2011



(B) Relative leisure travel attractiveness in 2015





Figure E.2: Relative leisure travel attractiveness in 2011 and 2015

Notes: These two figures show the spatial distribution of relative leisure travel attractiveness in 2011 and 2015, estimated based on the destination fixed effects from the leisure travel gravity estimations for 2011 and 2015, corresponding to Columns (2) and (6) of Table 3. These destination fixed effects are then calculated based on Equation (30) to yield each city's leisure travel attractiveness relative to the geometric average of all estimated destination fixed effects.

(A) Relative accessibility in 2011



(B) Relative accessibility in 2015




Figure E.3: Relative leisure travel accessibility in 2011 and 2015

Notes: These two figures show the spatial distribution of relative leisure travel accessibility in 2011 and 2015, estimated based on the origin fixed effects from the leisure travel gravity estimations for 2011 and 2015, corresponding to Columns (2) and (6) of Table 3. These origin fixed effects are then calculated based on Equation (31) to yield each city's leisure travel accessibility relative to the geometric average of all estimated origin fixed effects.




Figure E.4: Decomposition of the change in national average business travel time by year

Notes: This figure plots the year-by-year decomposition of the change in national average business travel time for each year from 2008 to 2015, relative to the no-HSR baseline, using Equation (40). Because the model is static, each bar is the long-run spatial equilibrium given that year's HSR network, holding baseline fundamentals fixed; the series is not a dynamic transition. Business travel is measured by procurement-related search volumes, with residual business travel scaled in proportion. Baseline shares are measured from 2011 UnionPay volumes.




Figure E.5: Business travel time reduction in 2015 by city

Notes: This figure maps the 2015 counterfactual percentage change in city-level average business travel time relative to the no-HSR baseline, according to Equation (D.1). Panel A reports the total change and Panel B reports the corresponding direct time effects. Business travel is measured by procurement-related search volumes, with residual business travel scaled in proportion. Hatched areas indicate cities outside the counterfactual sample.




Figure E.6: Leisure travel time reduction in 2015 by city

Notes: This figure maps the 2015 counterfactual percentage change in city-level average outbound leisure travel time relative to the no-HSR baseline, according to Equation (D.2). Panel A reports the total change and Panel B reports the corresponding direct time effects. Hatched areas indicate cities outside the counterfactual sample.




Figure E.7: Percentage change in resident population due to the 2015 HSR network by city

Notes: This figure maps the percentage change in each city's equilibrium resident population due to the 2015 HSR network relative to the no-HSR baseline. Blue denotes population gains and red denotes losses. Hatched areas indicate cities outside the counterfactual sample.




Figure E.8: Percentage change in the national welfare due to HSR by year

Notes: This figure reports the percentage change (calculated as 100 times the log change) in national welfare under each annual HSR network relative to the no-HSR baseline. National welfare is the *hukou*-population-weighted geometric mean of origin-specific expected utilities. “Both channels” allows HSR-induced travel-time reductions to affect both leisure travel and business search. The search-only and leisure-only counterfactuals hold the other friction at its baseline level. The synergy effect is the full effect minus the two channel effects.




Figure E.9: Percentage change in the real wage and *hukou*-residence welfare due to the 2015 HSR network by city

Notes: This figure maps the counterfactual percentage change in city-level real wage and *hukou*-residence welfare due to the 2015 HSR network relative to the no-HSR baseline. Panel A reports the real wage change and Panel B reports the *hukou*-residence welfare change. Hatched areas indicate cities outside the counterfactual sample.




Figure E.10: Counterfactuals with restricted HSR network reachability

Notes: This figure compares the baseline counterfactuals with a robustness exercise that shuts down HSR-induced travel-time reductions for city pairs more than 1,200 km apart in straight-line distance, a range at which air travel is a relevant substitute (Fang et al., 2025). For those pairs we set $\hat{t}_{nj} = 1$ and leave all other pairs and all estimated parameters unchanged. The first panel reports the percentage change in the national average real wage. The second panel reports the percentage change (calculated as 100 times the log change) in national welfare. In each panel, “All city pairs” is the baseline series and “City pairs within 1,200 km” is the restricted-shock series. All changes are relative to the no-HSR baseline.




Figure E.11: Direct HSR exposure, number of residents, and local real wages by city size

Notes: Panel A plots the model-implied relative change in resident population, $\hat{L}_n - 1$, and Panel B plots the relative change in the local real wage, $\hat{w}_n / \hat{P}_n - 1$, under the 2015 HSR network relative to the no-HSR baseline. The horizontal axis is the positive direct reduction in city-level average business travel time, calculated as the city-level counterpart of the direct time effect in Equation (40) while holding baseline travel shares fixed. The sample is restricted to cities with positive direct business-travel exposure. A city is large if its population in the no-HSR baseline exceeds the sample median. Dashed lines report group-specific linear fits.

Table E.1: Leisure travel gravity estimates under alternative distance thresholds

	Dependent variable: $\lambda_{m n,t}^T$				
	≤ 600 km	≤ 800 km	$\leq 1,000$ km	$\leq 1,200$ km	$\leq 1,400$ km
	(1)	(2)	(3)	(4)	(5)
$\ln t_{nmt}$	-2.184*** (0.080)	-2.178*** (0.073)	-2.146*** (0.076)	-2.085*** (0.084)	-2.045*** (0.085)
Dest. $\times$ year FEs	Yes	Yes	Yes	Yes	Yes
Origin $\times$ year FEs	Yes	Yes	Yes	Yes	Yes
# of observations	88,711	142,312	199,850	258,210	315,670

Notes: This table reports the PPML estimation results for Equation (29) under varying sample restrictions. The sample period spans 2011 to 2015, consistent with Column (1) in Table 3. To mitigate the potential overestimation of travel frictions for distant city pairs, the estimation samples are progressively restricted based on proximity. Columns (1) through (5) exclude city pairs with straight-line distances exceeding 600, 800, 1,000, 1,200, and 1,400 km, respectively. Standard errors, two-way clustered at the origin and destination city levels, are reported in parentheses.

*** p < 0.01, ** p < 0.05, * p < 0.1.

Table E.2: Structural parameters summary

Parameter	Description	Value	Source
<i>Panel A. Calibration</i>			
$\{\rho_n^k\}_{n \in N, k \in K}$	sector-specific final-good expenditure share in all final goods consumption	varies by city and sector	calibrated from the city-level MRIO table in 2012
$\{\gamma_n^k\}_{n \in N, k \in K}$	production elasticity of labor	varies by city and sector	calibrated from the city-level MRIO table in 2012
$\{\gamma_n^{k,s}\}_{n \in N, k,s \in K}$	production elasticity of materials	varies by city and sector	calibrated from the city-level MRIO table in 2012
ζ	share of capital cost in total housing construction cost	0.63	calibrated from Wu et al. (2012)
β	share of final goods consumption in total consumption	0.85	calibrated from <i>China Statistical Yearbook</i> for 2011
θ^R	elasticity of migration with respect to real wages	1.9	calibrated from Fang and Huang (2022)
θ^P	shape parameter of the idiosyncratic productivity distribution	3.42	solved from $\kappa\theta^P/(\kappa - 1) = 5.94$ which is calibrated from Che et al. (2024)
σ	elasticity of substitution between firm-specific final goods	6	calibrated from Anderson and van Wincoop (2004)
<i>Panel B. Estimation</i>			
$\alpha^T\theta^T$	elasticity of leisure travel with respect to travel time	2	estimated from the leisure travel gravity equation
$\theta^R\psi/\theta^T$	elasticity of migration with respect to leisure travel accessibility	0.18	estimated from the migration gravity equation
α^P	elasticity of business search volume with respect to travel time	0.3	estimated from the business volume regression
κ	convexity of search cost	2.36	estimated from the trade gravity equation

Notes: This table summarizes the structural parameters required for the counterfactual simulations, including their descriptions, values, and sources (either calibrated or estimated).

Table E.3: Sensitivity of the aggregate effects to the calibrated trade elasticity

Calibrated $\kappa\theta^P/(\kappa - 1)$	National real wage/GDP	Welfare
5.94 (baseline)	2.185	3.140
5.44	2.165	3.209
4.94	2.206	3.452
4.44	2.240	3.499
3.94	2.327	3.610

Notes: This table reports the effects of the 2015 HSR network relative to the no-HSR baseline under alternative calibrations of the composite trade elasticity with respect to freight costs, $\kappa\theta^P/(\kappa - 1)$. The baseline value of 5.94 is the average of the sector-specific estimates in Che et al. (2024). In each row, $\kappa = 2.36$ and all other parameters are held at their baseline values, so that θ^P equals 3.42, 3.13, 2.85, 2.56, and 2.27, respectively. All rows use the same no-HSR baseline allocation, with both the business-search and leisure-accessibility channels active. The percentage change in the national average real wage equals the change in national real GDP because national population is fixed. “Welfare” is 100 times the log change in national welfare. All calibrations satisfy the parametric restriction $\kappa\theta^P > \sigma - 1$.

Table E.4: Spatial heterogeneity in the effects of the 2015 HSR network by channel

	Large cities with HSR	Small cities with HSR	Cities without HSR
<i>Panel A. Search channel only</i>			
Local real GDP	2.014	1.137	−0.976
Resident population	0.247	0.041	−0.414
Real wage	1.762	1.096	−0.565
Welfare (<i>hukou</i> origin)	1.136	0.922	−0.391
<i>Panel B. Leisure channel only</i>			
Local real GDP	2.219	0.456	−1.680
Resident population	0.676	0.070	−1.119
Real wage	1.533	0.386	−0.567
Welfare (<i>hukou</i> origin)	4.231	3.495	−0.077

Notes: This table reports percentage changes under the 2015 HSR network relative to the no-HSR baseline. Panel A activates only the business-search channel, and Panel B activates only the leisure-accessibility channel. Outcome definitions, aggregation, and city classifications are the same as in Table 9.