

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Economics

School of Economics

3-2024

The distributional impacts of transportation networks in China

Lin MA

Singapore Management University, linma@smu.edu.sg

Yang TANG

Nanyang Technological University

Follow this and additional works at: https://ink.library.smu.edu.sg/soe_research



Part of the [Asian Studies Commons](#), [International Economics Commons](#), and the [Transportation Commons](#)

Citation

MA, Lin and TANG, Yang. The distributional impacts of transportation networks in China. (2024). *Journal of International Economics*. 148, 1-21.

Available at: https://ink.library.smu.edu.sg/soe_research/2721

This Journal Article is brought to you for free and open access by the School of Economics at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Economics by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

The Distributional Impacts of Transportation Networks in China*

Lin Ma Yang Tang[†]

December 2023

Abstract

We document that the quality of roads and railroads vary substantially over time and space in China, and neglecting these variations biases the distributional impacts of transportation networks. To account for quality differences, we construct a new panel dataset and approximate quality using the design speed of roads and railroads that varies by vintage, class, and terrain at the pixel level. We then build a dynamic spatial general equilibrium model for multiple modes, transportation routes, and forward-looking migration decisions. Our findings demonstrate that disregarding quality differences leads to a median bias of approximately 31% in estimating real wage growth rates at the prefecture level. Moreover, this bias is non-random and correlates with the initial conditions of the prefectures, resulting in significant errors when predicting the distributional effects of transportation networks.

Keywords: regional trade; migration; welfare; economic geography

JEL Classification: F1; F4; R1; O4

*We are especially grateful for valuable insights from the editor, two anonymous referees, and stimulating discussions with Davin Chor, Wen-Tai Hsu, Stephen Redding and Daniel Xu. We have also greatly benefited from comments at various conferences and seminars, including the Asian Bureau of Finance and Economic Research, Asia-Pacific Trade Seminars, and the Federal Reserve Bank of St.Louis. Yang Tang acknowledges financial support from Singapore MOE-Tier 1 research fund(M4011853).

[†]Singapore Management University (linma@smu.edu.sg) and Nanyang Technological University (tangyang@ntu.edu.sg), respectively.

1 Introduction

Over the past decades, developing countries worldwide have made enormous investments in transportation networks.¹ While it is well understood that the improvements in transportation infrastructure bring economic prosperity on average, the distributional impacts across space are controversial: it is unclear if better connectivity widens or narrows income gaps across locations.² From a theoretical perspective, the predictions are ambiguous: on the one hand, the reductions in trade frictions allow remote locations to gain better market access and reduce spatial inequality. On the other hand, better connectivity also improves factor mobility that drains resources away from remote locations.

In addition to the theoretical ambiguity, the variation in the quality of the transportation infrastructure over time and space poses another significant challenge. The poorer and more isolated towns often reside in rugged terrains, making construction challenging. As a result, the roads in these areas may have lower capacity and design speed, even though they may share the same designation as those in well-developed regions. At the same time, the remote locations are connected to the transportation network in later periods when a country’s engineering capacity and technology have considerably improved, leading to cross-time differences in infrastructure quality. Measuring the quality of infrastructure is thus crucial to capture the distributional effects of the transportation network.

This paper provides the first dataset on roads and railroads that explicitly accounts for the quality variations across space and time for China — a country that has experienced a spectacular expansion in its transportation network over the past 40 years. In this dataset, we go beyond the binary connectivity measures in the literature and measure the *quality* of road and railroad at the pixel level. Existing datasets often geo-reference the highways and rely on binary indicators of whether a road exists on a pixel to measure connectivity.³ The implicit assumption behind

¹For example, [Gurara et al. \(2018\)](#) estimated that developing countries, on average, spend around 6 percent of GDP on infrastructure each year. The World Bank estimates that in the next decade, developing countries need to spend around 3.3 percent of GDP on transportation networks to meet the future demands for mobility, according to [Rozenberg and Fay \(2019\)](#).

²For example, in the context of China, [Faber \(2014\)](#) and [Baum-Snow et al. \(2020\)](#) document that highway connectivity hurts the development of small and less connected locations, while [Egger et al. \(2023\)](#) show that better connectivity improves spatial convergence of real income. See [Redding and Turner \(2015\)](#) for a review of the underlying theoretical issues.

³For example, see the road connectivity measures in [Banerjee et al. \(2012\)](#), [Faber \(2014\)](#). The data in [Baum-Snow et al. \(2020\)](#) and [Egger et al. \(2023\)](#) differentiate quality by class of roads, such as highways and national roads. However, these data assume no quality variation across time and space within each class. The maps on the Chinese road system hosted by the China Data Center at the University of Michigan, the United States Geological Survey (USGS), the ACASIAN Data Center at Griffith University, the OpenStreetMap Project, and the State Bureau of

these exercises is that all the highways are built with the same quality. However, this is not true. The official engineering standards for roads and railroad design have improved substantially over the years. As a result, traffic on the roads constructed in the earlier years cannot travel as fast as on the recent ones. Moreover, conditional on the vintage and rate of the road, quality also varies across space. Due to engineering difficulties, the road and railroads' design speed in the mountainous regions could be only half of those in the plain regions. Lastly, the spatial and temporal variations in quality also interact. As the highways and railroads constructed in recent years were mainly placed in the less developed regions with rugged terrains, the spatial variations in quality could also lead to biased estimates of cross-time differences in transportation networks. In the empirical literature, the inconsistency of quality standards over time is widely considered a significant obstacle to exploring the time-series variations in transportation networks in China. For example, [Baum-Snow et al. \(2020\)](#) observed that

“...however, the growth and improvement of China’s road network were so dramatic that roads that were important enough to merit inclusion on the 1990 map probably bear little resemblance to roads that meet this standard in 2010, even if both roads received the same designation in the legend. Thus, we are reluctant to exploit the time-series variation in our measures of highways.”

To account for the variations in infrastructure quality, we document the revisions in design codes over several decades using official publications from the Ministry of Transportation in China. Each revision of the design code stipulates the design speed of roads and railroads by the rate and terrain: higher-rated roads are often designed to allow for faster travel speed as compared to lower-rated ones, and rugged terrains such as hills and mountains usually impose stricter limits on the design speed. Specifically, we first break each transportation network into segments. For each segment, we then identify the year of construction and the rate of the road and railroad in each year by cross-referencing a wide array of sources, including the *Transportation Yearbooks*, provincial map collections, and the *Chronicles of Railroad Construction* ([Ma, 1983](#)). Lastly, we determine the terrain type at the pixel level for each segment and apply the respective design code to each pixel. These steps combined allow us to determine the “design speed,” which we interpret as the quality at the pixel level for every road and railroad in China. The resulting dataset allows us to measure the distance between two pixels by travel time. To our knowledge, the accomplished dataset is the Surveying and Mapping of China are all binary maps without any measure of quality.

first to measure the transportation network in a country consistently over time and across space, and indeed the first for China.

Besides the contribution in quality measures, our dataset surpasses existing data in three key aspects. Firstly, our dataset spans a significantly longer period than previous works. For instance, the railroad component of our dataset encompasses every railroad constructed in China up until 2017, starting from the very first railroad built during the late Qing Dynasty in 1881. Similarly, the highway component encompasses every highway in China from 1988 to 2017. This comprehensive coverage of railroads and highways enables researchers to trace the network’s development year by year.⁴ Secondly, due to variations in projection methods employed by different map publishers, the geo-referencing of the same road segment can be inconsistent across different years when the source maps originate from diverse publishers. To tackle this issue, we have implemented an iterative procedure during the dataset construction, ensuring that every pixel of a road maintains consistent coordinates throughout the years. Lastly, we distinguish between freight and passenger travel within the railroad networks. This differentiation is particularly important in China, as the High-Speed Rail (HSR) system exclusively caters to passenger traffic. This distinction delivers significantly different welfare implications as the freight networks directly impact internal trade but not factor mobility, and vice versa for the passenger networks.

Our dataset reveals that the expansions in transportation infrastructure from 1994 to 2017 have significantly reduced the average travel time. For the road network, the median value of the reduction in travel time among all city pairs amounts to 32 percent. For the railroad network, the median reduction reaches 37 percent for freight travel and 59 percent for passenger travel. More importantly, ignoring the variations in quality can lead to substantial bias in measuring the reduction in travel time. In the case of the road network, assuming uniform quality would result in an *overestimate* of the median reduction by 22 percent. This bias arises from the fact that newly-constructed highways are often located in mountainous regions with low design speeds. On the contrary, the absence of quality measures would *underestimate* the improvements in median travel time on the rail system by 46 percent for freight travel and 61 percent for passenger travel. By assuming uniform quality, the substantial advancements in railway design speeds over time are disregarded, resulting in an underestimation.

The quality-adjusted travel time estimates from our data can better explain the variations in

⁴As explained in detail later, we also cover the first-rated roads, commonly known as “national roads” in the literature. Unlike highways and railroads, we take the existing network in 1994 as given and only trace the development from 1994 to 2017 annually.

observed travel time on both the road and the rail network than the unadjusted ones. To do so, we compare the travel time estimations with and without quality adjustment to the observed timetable data from long-distance bus and passenger rail service providers. On the road system, both travel time estimates highly correlate with the recorded travel time of the long-distance buses. However, when we control for both travel time measures and run a horse race between them, the estimate without quality adjustment is no longer significant. In contrast, the coefficient on the quality-adjusted travel time is barely affected. Results are similar on the railroad system. As we do not use any timetable information in compiling the dataset, these results indicate that the quality adjustments in our dataset capture the actual variations in travel speeds on the road and rail network.

To evaluate the general equilibrium impacts of transportation networks, we develop a dynamic spatial general equilibrium model in which trade and migration flows are subject to bilateral frictions that depend on infrastructure. We model route choices in the transportation network following [Allen and Arkolakis \(2022\)](#) and extend their framework to allow for multiple modes of transportation. We embed the mode and route choice problem into a dynamic discrete choice framework as in [Caliendo et al. \(2019\)](#) and allow the migration decisions to be forward-looking. In the model, individuals' location choices not only depend on the current state of infrastructure but also the anticipated changes in the future. In addition, we also include China-specific elements, such as the policy barriers to migration in the form of the hukou system.

Counterfactual simulations show that expanding transportation networks improved aggregate welfare and reduced spatial inequality. The expansion of transportation networks between 1995 and 2017 improves aggregate welfare by 57 percent in the long run. The aggregate gain comes from lowered internal trade frictions and better population allocation across space. Moreover, the expansion of transportation networks significantly reduces spatial inequality. In the counterfactual simulation without infrastructure improvements, spatial inequality barely changes. In stark contrast, real wage exhibits regional convergence under various inequality measures in the baseline simulation that incorporates the actual changes in transportation networks. Furthermore, we find almost all the convergence in economic activities during that period comes from the changes in transportation networks.

More importantly, disregarding the quality variations leads to considerable bias in model predictions. When comparing the baseline simulation with a counterfactual scenario assuming uniform quality of roads and railroads across time and space, we observe substantial measurement error in

real wage growth rates at the prefecture level. Along the transition path, the median error ranges from 19% to 33%. At the 99th percentile of the relative error distribution, the bias spans an extensive range from 175% to 542%. In other words, neglecting quality variations can result in an almost five-fold overestimation of growth in real wage in certain prefectures.

The measurement errors that come from ignoring quality are not random; instead, they correlate with the initial conditions of the prefecture. This non-random nature suggests two important implications. Firstly, empirical researchers should not consider these errors as classical measurement errors. As a result, typical solutions such as instrumental variables may not effectively address the measurement errors arising from quality variations. Second, the correlation also implies that the distributional impacts of transportation networks are particularly sensitive to variations in quality. We show that assuming uniform quality overstates the importance of transportation networks in reducing spatial inequality, and the bias ranges between 19% and 36% depending on timing and inequality measures.

Our paper contributes to a large literature on the economic impacts of infrastructure improvements (Baum-Snow, 2007; Banerjee et al., 2012; Faber, 2014; Redding and Turner, 2015; Donaldson and Hornbeck, 2016; Qin, 2016; Lin, 2017; Xu, 2017; Donaldson, 2018; Baum-Snow et al., 2020; Fan et al., 2021; Alder et al., 2023; Bonadio, 2022; Wong and Fuchs, 2022; Egger et al., 2023). In the context of China, we provide the first dataset to measure the quality differences of infrastructure across time and space. The closest to our work is Egger et al. (2023) which documents the growth of transportation networks in China between 2000 and 2013. Their dataset categorizes roads and railroads by classes such as highways, national roads, and lower-tier roads. However, quality does not vary within each class. In addition to the longer time span in our data, we show that within each category, quality varies across time and space, and ignoring such differences could lead to considerable biases in model predictions. We also add to a small and growing literature highlighting the economic importance of infrastructure quality (Berg et al., 2018; Asher and Novosad, 2020; Jedwab and Storeygard, 2021). We go beyond the categorical measure of quality in this literature (paved v.s. non-paved roads, highway v.s. non-highway) and instead measure quality using design speed that changes over time and space, and can be directly mapped into travel time costs.

Our work also contributes to the literature of quantitative spatial models (Allen and Arkolakis, 2014; Ahlfeldt et al., 2015; Redding, 2016; Tombe and Zhu, 2019; Caliendo et al., 2019; Kleinman et al., 2023; Allen and Arkolakis, 2022). The closest to our model are Allen and Arkolakis (2022) and Caliendo et al. (2019). Relative to Allen and Arkolakis (2022), we allow for multiple modes of

travel and dynamic response of migration. Relative to [Caliendo et al. \(2019\)](#), we introduce route choices in the dynamic migration framework. Despite the significant changes in the choice sets, we show that the migration decision still adopts a tractable solution.

The rest of the paper is organized as follows. Section 2 introduces the panel dataset on the transportation networks of China; Section 3 presents the model and Section 4 the quantification. Section 5 discusses the results, and Section 6 concludes.

2 The Transportation Networks of China

In this section, we outline the construction of the panel dataset on transportation networks in China and refer the readers to Appendix B for details. We briefly discuss the underlying inconsistency in measuring the quality of roads and railroads across time and space. We then describe the basic patterns of the evolution of the transportation networks over this period in China and the importance of accounting for quality variations.

The temporal coverage of our dataset varies depending on the mode of transportation. For railroads and highways, we track the evolution of every individual railroad (1881 to 2017) and highway (1998 to 2017) from their inception in the country. We limit our analysis to 1994 to 2017 for the lower-tiered roads known as "first-rated roads." This is due to the lack of clear documentation regarding their construction in yearbooks and chronicles.⁵

Measuring Binary Connectivity The starting point of compiling the dataset is to collect the published transportation atlas for each year. We source the physical maps through several channels, such as libraries, used book dealers, and map collectors. We only choose the national-level atlas with a scale greater than 1:6 million for geo-referencing, as smaller maps do not provide enough resolution for color identification. The dimension of the digitized national maps is 12669 pixels in width and 8829 pixels in height, so each pixel is around 500 meters in length. Out of the 24 years, we have obtained national maps for 10 years. To identify the year of the construction of every road segment and fill in the gap years, we cross-referenced the annual *Collection of Provincial Transportation Maps*, transportation and railroad yearbooks, as well as published chronicles on transportation construction in China. Table A.1 lists the physical maps and other descriptive

⁵Alternatively, one could measure travel time using web-based services such as Google or Baidu Maps. These methods offer the advantage of considering real-time traffic compared to our dataset. However, our dataset has a distinct advantage in historical coverage. For example, none of these services provided historical transportation network data before 2005, while we traced the development of railroad and highway networks back to their inception.

references in our collection. At the end of this stage, our dataset identifies the construction year of each geo-referenced segment, from which one can trace the evolution of the connectivity by modes of transportation. Figure A.1 in the Appendix maps the evolution of the connectivity networks in several years.

2.1 Measuring Quality

Quality variations occur both over time and across space. With the rapid improvement in the capacity and capability of the civil engineering sector, the Chinese government has substantially lifted the highway and railway construction standards over the years. As a result, a “highway” constructed in 1988 might only be appropriately classified as a “first-rate road” by the 2014 standards. In addition, quality also varies across space. For example, under the 1988 construction standard, the design speed for highways in the eastern flood plains is 120 km/h and 60 km/h for those in the mountainous regions: assuming a uniform quality would significantly overestimate the improvements in connectivity if the newly constructed highways concentrate in mountainous regions.

We measure the quality of roads and railroads using their “*design speed*”. The design speed is an ideal choice for two reasons. First, other roads and railroad engineering standards correlate with the design speed. For example, a road with a higher design speed is usually built with more lanes, straighter distances, and stronger supporting materials. Therefore, this single variable sufficiently reflects the quality of a road. Secondly, from the design speed, one can directly compute the time cost of trespassing a pixel and, subsequently, the time cost of traveling between any two pixels on the map. The time cost of traveling can then be readily incorporated into a quantitative trade and migration model, as discussed in Section 3. In this section, we briefly discuss road and railroad quality measurement and refer the readers to Appendix B.2 for more details.⁶

Roads Panel (a) of Table 1 summarizes the design speed of roads by the year and terrain of the construction. These stipulated design speeds are based on the *Technical Standard of Highway Engineering*, published by the Ministry of Transportation of China (MOT).⁷ Nine revisions existed

⁶We also need to measure terrains in China in order to infer design speed. The definition of the terrains comes from the *Land Regulations in Highway Engineering*, published by the Ministry of Transportation, and the United Nations definition of terrains from Kapos et al. (2000).

⁷Despite the name of “highway” in the title, the standards regulate all inter-city roads, including the highways (Gao Su Gong Lu) and normal roads (Yi Ban Gong Lu) from the first-rate to the fourth-rate. “Highway” in the title is the official translation of “Gong Lu”, which means inter-city roads. These standards do not regulate the “intra-city roads” (Cheng Shi Dao Lu).

from 1951 to 2014, of which four are relevant for our sampling period: 1988, 1997, 2003, and 2014. Each revision defines the design speed for roads of different rates (highways, first, second, and third-rated roads). We only focus on the highways and the first-rated roads identified in the previous step. With the information on the year of the construction and the underlying terrain, we can determine the design speed for each road pixel.

Most quality variation in the road network comes from the spatial dimension. In the earlier standards, the design speed of all roads depends on the terrain. Over the years, the Ministry of Transport has gradually reduced the reliance on terrain for highways while still maintaining the terrain dependency for lower-rate roads.⁸

Railroads Panels (b) of Table 1 summarize the design speeds of mixed-use railroads by year and construction terrain.⁹ Similar to the standards in highway engineering, the railroad designs depend on the intended usage, the rate of a railroad (National I-IV, Industrial I-III), and the underlying terrain. In this section, we highlight several issues specific to measuring railroad quality.

First, the quality variation in the railroad system comes from both the temporal and the spatial dimensions. Similar to the road system, railroads in the hills and mountains limit travel speed. However, unlike the road system, the railroads' design speed changes significantly over time.

Second, the operation speed of existing railroads also changes over time due to engine and route design improvements. The Ministry of Railroads implemented six waves of speed improvements on the existing rail network, which increased the speed of old networks by as much as 30 percent in some cases. We incorporate such improvements in the dataset by referencing the yearbooks.

Third, railroad designs vary by intended usage and fall into three categories: passenger-only, freight-only, and mixed-use. The passenger-only railroads, including the High-Speed-Rail (HSR) system, emphasize travel speed but cannot handle heavy loads. The freight-only railroads prioritize load capacity over speed. Most of the railroads fall into the last category that lies in the middle: the mixed-use railroads. We refer to the yearbooks and the chronicles of railroad construction to identify the usage of each railroad.

⁸For instance, in the 2003 revision, it was mandated that all highways should be designed with a speed limit of 120 km/h, except in particularly challenging regions. Consequently, more elevated highways were constructed in low-rolling hills and hilly terrains.

⁹These tables are based on various publications from the Ministry of Transportation and the Ministry of Railroads, such as the *Code for Design of Railway Line*, *Code for Design of Standard Railway Line for Industrial Firms*, and the *Code for Design of III and IV-rated Railway Line*. See Table A.3 for the full list of design speed of freight railways.

Table 1: Design Speed (km/h) of Roads and Railroads by Time and Terrain

(a) Road Standards

Revision	Plains	LRH	Hills	Mountains	Plains	LRH	Hills	Mountains
	Highways				First-Rate Roads			
1988	120	120	100	60	100	100	60	60
1997	120	120	120	60	100	100	60	60
2003	120	120	120	80	100	100	80	80
2014	120	120	120	80	100	100	80	80

(b) Railroad Standards, Mixed-Use

Revision	Plains	LRH	Hills	Mountains	Plains	LRH	Hills	Mountains
	National I				National II			
1985	120	100	80	80	100	80	80	80
1999	140	120	80	80	100	80	80	80
2006	160	140	120	120	120	100	80	80
Revision	National III				National IV			
	Plains	LRH	Hills	Mountains	Plains	LRH	Hills	Mountains
1985	80	80	80	80	-	-	-	-
1999	80	80	80	80	-	-	-	-
2012	120	100	80	80	100	80	60	40

Notes: This table summarizes the design speed of the highways and the first-rate roads by revisions of the *Technical Standard of Highway Engineering*, and that of the railroads by various revisions of railroad engineering. “LRH” refers to “Low Rolling Hills”. Please refer to Appendix B.1 for more details.

Waterway To ensure consistency across the modes of transportation, we also measure the quality of waterway transportation by speed. As there are no significant improvements in sailing speed over the sampling period, we assign a sailing speed of 22.8 km/h to all the waterway pixels in all the years.¹⁰

Structure of the Dataset Our end product is a panel dataset that documents the evolution of design speeds each year at the pixel level by modes of transportation. Figure 1 illustrates the dataset’s structure using a hypothetical railway as an example.

The outermost layer of the dataset is a “path”, a group of segments that form a known road or railroad. For example, a “path” could refer to a group of segments that form the Beijing-Shanghai Highway. In addition to names, paths vary in “rate” and “usage”. In the road dataset, the rate of a path could be “Highway” or “First-Rated Roads”. In the railroad dataset, a path can be classified as “National I, II,...”. We use the rate of the roads and railroads to determine the applicable design codes at the segment level. The railroad usage also varies at the path level, which could be

¹⁰This estimate is the median sailing speed based on a sample of 5,112 vessels between 2012 and 2018 extracted from the Automatic Identification System (AIS) database. The median sailing speed only varies between 20.7km/h and 23.2km/h across years in our sample.

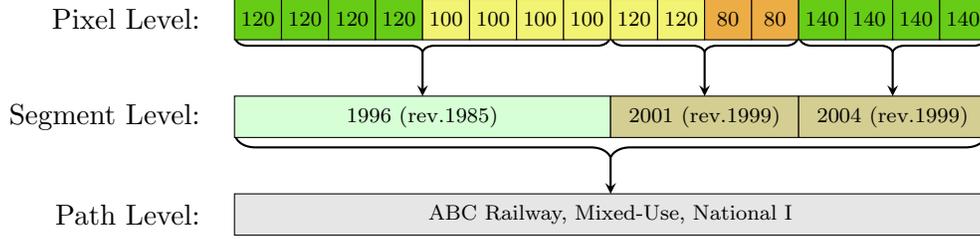


Figure 1: Structure of the Dataset

Notes: This diagram explains the dataset’s structure using a hypothetical railway, “ABC Railway”, as an example. At the path level, we identify the usage (mixed-use), rate (National I), and the path’s name. This path contains three segments, built in 1996, 2001, and 2004. The segment built in 1996 was subject to the 1985 revision of *Code of Design of Railway Line*, while the two segments built in 2001 and 2004 were subject to the 1999 revision. Each segment contains pixels that differ by terrain. ■ indicates “plains”, ■ indicates “low-rolling hills (LRH)”, and ■ indicates “hills”. The number in the pixel boxes is the design speed at each pixel in the unit of “km/h”. Within a segment, design speed differs due to terrain; within a terrain type, design speed differs due to the changes in applicable technical standards.

passenger-only, freight-only, or mixed-use. All the road paths are mixed-use.

One layer below the “path” is the “segment”. Within a path, segments vary in the year of the construction. It is common for a path to be constructed over many years, and we record these variations at the segment level. The year of construction then determines the applicable revision of the technical standards. As an example shown in Figure 1, the segment constructed in 1996 was subject to the 1985 revision of design codes, while the two segments built in 2001 and 2004 were subject to the 1999 revision.

At the lowest level of the dataset are the “pixels” that form each segment. Within a segment, pixels differ in their underlying terrain. Together with the rate, usage, and year of the construction from the upper layers of the dataset, we can then determine the design speed of each pixel. In Figure 1, the pixel’s color indicates the terrain, and the number in each pixel states the design speed in the unit of “km/h”. Note that within a segment, design speed differs due to terrain; within a terrain type, design speed differs due to the changes in applicable technical standards.

2.2 The Evolution of Transportation Networks in China

Based on the design speed at the pixel level identified in the previous part, we compute the travel time between two prefectures, i, j , in year t , using the fast-marching algorithm. We denote the time cost matrix in the unit of hours as $\mathbf{T}_t^{m\chi}$, where the (i, j) th element of the matrix is the travel time from j to i under transportation mode m for traffic type χ . The mode of transportation, m , takes three values: $m = 1$ indicates the road, which includes both the first-rate roads and the highways;

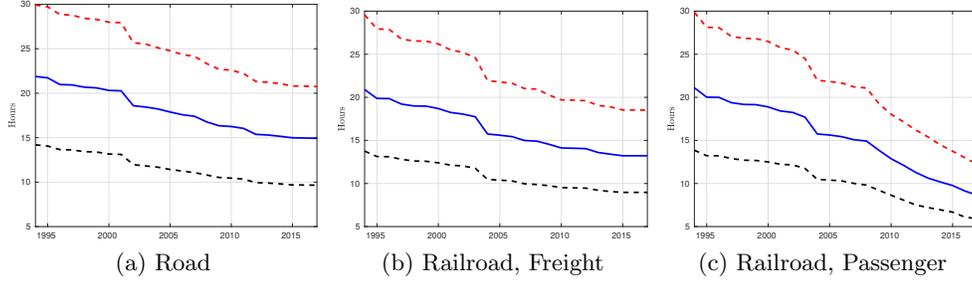


Figure 2: The Evolution of Travel Time

Notes: The figures present the changes in travel time across all prefecture pairs by modes of transportation and types of traffic. In all the panels, the solid blue line in the middle is the median, the dashed red line is the 75th, and the dashed black line is the 25th percentile of the travel time distribution across all the prefecture pairs. The travel time on the road network is the same for freight and passenger traffic; therefore, we only present one set of figures under the sub-caption “road”. The figure for waterway transportation is omitted because it does not vary over time.

$m = 2$ is the railroad, and $m = 3$ is the waterway network. We compute the travel time separately for two types of traffic denoted by χ , in which $\chi = \text{f}$ indicates the freight traffic, and $\chi = \text{p}$ is the passenger traffic. Both traffic types travel on the road and waterway networks. On the railroad network, freight traffic cannot utilize the passenger railroads, and the passenger traffic cannot use the freight railroads. Figures 2 to 5 present the basic findings. Several messages are conveyed as follows.

The first message is that expanding transportation networks significantly reduced travel time in China. As shown in Figure 2, the median travel time between all prefecture pairs on the road networks declined by around 32 percent from nearly 22 hours in 1994 to approximately 14 hours in 2017. The reduction in the travel time on the rail networks is more substantial, at 37 percent for freight travel and 59 percent for passenger travel. The significant reduction in passenger travel time is due to the expansion of the passenger-only HSR network. Moreover, the improvements in connectivity occurred across the board: at the 75th percentile of the travel-time distribution, we observe a reduction of similar magnitude, and at the 25th percentile, a slightly smaller reduction. The decreasing marginal returns in connectivity are intuitive. While a road connecting two remote prefectures reduces travel frictions substantially, the same might not be valid for a prefecture-pair closer to the center of the network.

The second message is that the reduction in travel time varies substantially across prefecture pairs, modes, and time. Figure 3 showcases this pattern by plotting the normalized travel time from various origins to Beijing. Some city pairs, such as Beijing-Tianjin in the first panel, experienced a negligible reduction in travel time on the road network, and others, such as Beijing-Wulumuqi, saw

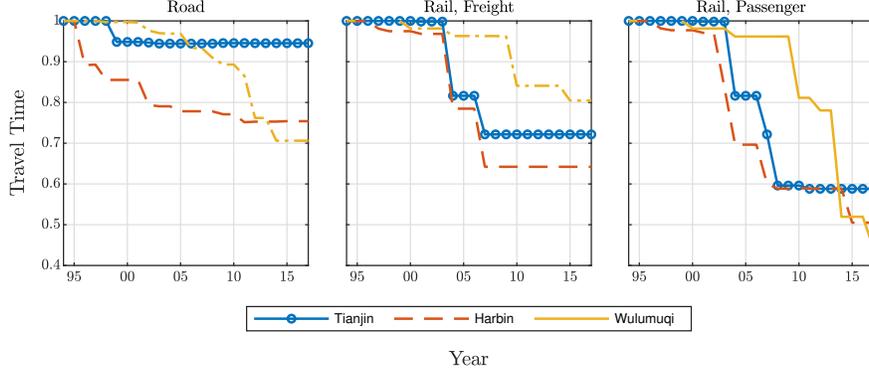


Figure 3: The Evolution of Transportation Networks, Distance to Beijing

Notes: The figures present the travel time between Beijing and three selected destinations over time. The travel time is normalized to 1 in the initial year. The figure for waterway transportation is omitted because it does not vary over time.

a decline of close to 30 percent. Along the temporal dimension, changes in travel costs also vary. For example, while the Beijing-Haerbin road connection improved before 2000, those between Beijing and Wulumuqi only improved in 2012 after the construction of the Lanzhou-Wulumuqi Highway. Figure A.2 in the appendix summarizes the travel time reduction towards other prefectures, and the messages are broadly consistent with those presented in Figure 3.

The third message is that those initially remote cities received a more considerable reduction in travel time. Figure 4 plots the reduction in the average travel time between 1994 and 2017 from an origin to all other prefectures against the initial average travel time in 1994. A positive relationship emerges in the rail network transportation, and a moderate positive relationship can also be seen in the road network. The uneven improvement in the connectivity is due to the asymmetry in the initial network layout, in which the coastal cities were already well-connected to each other while the infrastructure in the inland regions was sparse in 1995. Under this circumstance, linking a remote city to the extensive network along the coast will reduce the travel time of the inland city to all the coastal cities in the existing network, substantially reducing its average travel time. In the other direction, the additional access to the remote city barely affects the average connectivity of the coastal cities.

The results also hold at the prefecture-pair level. Table 2 regresses the changes in travel time at the prefecture-pair level to various pair characteristics. Relative to the reference group, having one remote prefecture in the pair significantly amplifies the reduction in travel time.¹¹ In addition,

¹¹A “remote” location refers to a prefecture with above-median average travel time to all the other prefectures in 1994. A “central” location has below-median average travel time. The reference group is the pair in which both the origin and the destination prefecture are centrally located on the transportation network.

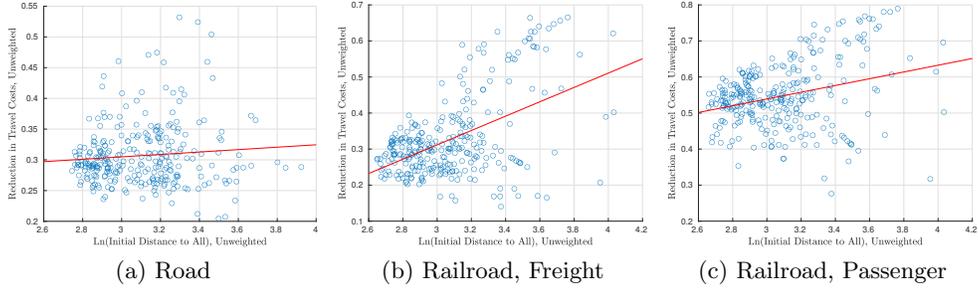


Figure 4: The Reduction in Transportation Costs v.s. Initial Position

Notes: This figure plots the reduction in travel time between 1994 and 2017 against the (logarithm of) the initial average travel time of each prefecture to all the other prefectures in 1994. Each dot represents a prefecture. The reduction in travel time is measured as $1 - T_{ij,2017}^m / T_{ij,1994}^m$, where T_{ijt}^m is the travel time between prefectures i and j in year t via mode m . The red line is the best linear fit. The figure for waterway transportation is omitted because it does not vary over time.

if both locations are remote, the reduction in travel time is even more pronounced, as illustrated in the first two rows of Table 2.

In addition to the initial location, the improvements in connectivity are more substantial in locations with lower initial population or GDP. Prefecture pairs with a higher initial (average) population or initial (average) GDP usually see smaller reductions in travel time, as suggested by the last two rows from Table 2. The heterogeneity in prefecture size is expected, as remote prefectures are often smaller in both population and output than the central ones.

Table 2: Reduction in Travel Time and Initial Conditions

	Road		Railroad, Freight		Railroad, Passenger	
	(1)	(2)	(3)	(4)	(5)	(6)
One Remote	0.005*** (0.001)	0.006*** (0.001)	0.082*** (0.002)	0.085*** (0.002)	0.076*** (0.002)	0.080*** (0.002)
Both Remote	0.008*** (0.001)	0.009*** (0.001)	0.134*** (0.002)	0.142*** (0.002)	0.127*** (0.002)	0.136*** (0.002)
Ln(ini. Pop)	-0.008*** (0.000)		-0.039*** (0.001)		-0.040*** (0.001)	
Ln(ini. GDP)		-0.008*** (0.000)		-0.012*** (0.001)		-0.010*** (0.001)
N	39,340	39,340	39,340	39,340	39,340	39,340
Adj.R2	0.009	0.011	0.101	0.078	0.099	0.074

Notes: This table reports the regressions of the reduction in transportation costs against various measures of initial conditions. Robust standard errors are in parentheses. ***: significant at the 1% level; *: significant at the 5% level; *: significant at the 10% level. Each observation is a prefecture pair. The dependent variable is $1 - T_{ij,2017}^m / T_{ij,1994}^m$, where T_{ijt}^m is the travel time between prefectures i and j in year t via mode m . “One Remote” indicates either i or j (but not both) has an average initial travel time above the median, and “Both Remote” indicates that both i and j ’s initial travel times are above the median. The reference group is “None Remote”: both i and j have below-median initial travel time. “Ln(ini. Pop)” and “Ln(ini. GDP)” are the logarithm of the average initial population and GDP of i and j , respectively.

2.3 The Importance of Quality

Before assessing the significance of quality variations using a structural model, we demonstrate the importance of accounting for quality with several reduced-form exercises. We first re-calculate the between-prefecture travel times, assuming uniform quality for roads and railroads. Specifically, we compute three sets of “uniform quality” travel times. In the first set, named “uniform high quality,” we assume all the infrastructure was constructed using the highest possible design speed, corresponding to those on plains in the latest engineering standards. The second set, named “uniform low quality,” adopts the lowest possible design speed, corresponding to those on mountain terrains in the earliest engineering standards. The last set, named “uniform quality”, seeks a middle ground and uses the median design speed of those on LRH and hills between the earliest and the latest engineering standards.¹² In the rest of this section, we focus on the last group, “uniform quality,” and compare the results based on our baseline measures with quality adjustments to those in the uniform measures without variations in quality. We refer the readers to Figure A.3 and Tables A.5 and A.6 for the results related to the “uniform high (low) quality” measures.

Changes in Travel Time Figure 5 presents the ratio of between-prefecture travel times in 2017 compared to 1994, both in the baseline setup considering quality differences and the setup assuming uniform quality. Panel (a) depicts the improvements in road connectivity, and we observe that the absence of quality differences would have led to an *overestimation* of these improvements. While the baseline case shows a median reduction in road travel time of 32 percent, the uniform quality case demonstrates a significantly higher reduction of 39 percent. This *overestimation* of approximately 22 percent can be attributed to the fact that newly constructed highways are often located in mountainous regions of China, where low design speeds are required.

In contrast, assuming uniform quality for railroad segments would have *underestimated* the improvements in railroad connectivity, as depicted in panels (b) and (c) of the same figure. For instance, the baseline case shows a median reduction of 37 percent for freight transportation and 59 percent for passenger transportation. However, under the uniform quality case, the reduction in travel time is only 20 percent for freight and 23 percent for passenger travel, resulting in underestimations of 46 percent and 61 percent, respectively.¹³ Unlike highway design codes that exhibit limited temporal variation, the design speed of railroads improved significantly over time.

¹²Table A.4 in the appendix provides the corresponding design speeds used in these uniform quality measures.

¹³Calculations: $1 - 20/37 \approx 0.46$, and $1 - 23/59 \approx 0.61$.

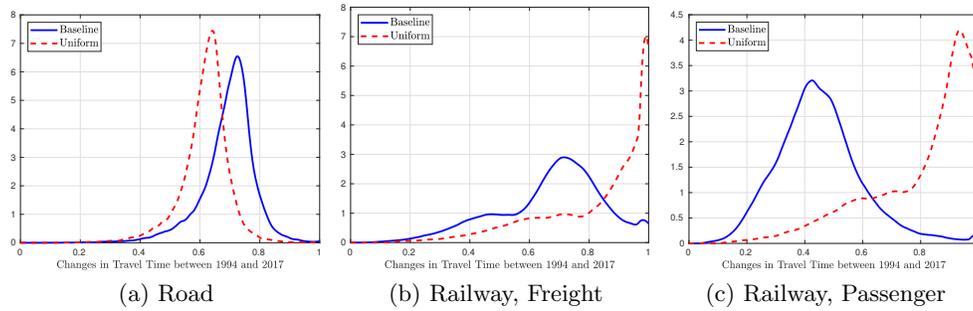


Figure 5: Comparison between Baseline and Uniform Quality

Notes: The figures present the kernel density plot of the ratio of travel time between all prefectures in the year 2017 to those in 1994. The blue line labeled “baseline” accounts for the variations in quality as described in the main text. The red line labeled “uniform” assumes uniform quality across time and space. In this case, the design speed for all the highway pixels is 120 km/h, for first-rate roads, 80km/h, and for all the railways, 100km/h.

Neglecting these temporal variations would result in underestimations.

Correlation with Initial Conditions The measurement errors introduced by ignoring quality adjustments are not random. Instead, they correlate with the initial conditions of the locations. To highlight this point, we first measure the bias introduced by ignoring quality as $\left|T_{ij,\text{uniform}}^m/T_{ij,\text{baseline}}^m - 1\right|$, where $T_{ij,(\cdot)}^m$ is the travel time between prefectures i and j in 2017 via mode m under the quality-adjusted “baseline” and the unadjusted “uniform” measures of infrastructure. We then regress these measurement errors against the same regressors listed in Table 2 and report the results in Table 3.

The measured errors between the baseline and the “uniform quality” scenario tend to be larger for initially remote, less populated, and poorer locations on the road network. As indicated in the first two columns of Table 3, relative to the reference group in which both locations are central, having a remote location in the pair increases the measurement errors by around 2.4%; when both locations are remote, the measurement errors grow to between 3.8% to 4.1%. Locations-pairs with lower average population and GDP also tend to see larger measurement errors. The measurement errors are also highly correlated with the initial conditions on the rail network, though some opposite patterns from the road network appear. For example, the measurement errors on the passenger rail network are *higher* among more centrally-located prefecture pairs. This is likely because the HSR system and six waves of speed adjustments mainly improved the network connecting major population centers but not the peripheral ones. We will later show in Section 5 that these measurement errors in travel time lead to substantial biases in predicted growth rates

Table 3: Measurement Errors and Initial Conditions

	Road		Railroad, Freight		Railroad, Passenger	
	(1)	(2)	(3)	(4)	(5)	(6)
One Remote	0.024*** (0.001)	0.025*** (0.001)	0.038*** (0.002)	0.035*** (0.002)	-0.054*** (0.006)	-0.055*** (0.006)
Both Remote	0.038*** (0.001)	0.041*** (0.001)	0.043*** (0.003)	0.036*** (0.003)	-0.174*** (0.007)	-0.179*** (0.006)
Ln(ini. Pop)	-0.014*** (0.000)		0.076*** (0.001)		0.158*** (0.003)	
Ln(ini. GDP)		-0.007*** (0.000)		0.062*** (0.001)		0.151*** (0.002)
N	39,340	39,340	39,340	39,340	39,340	39,340
Adj.R2	0.123	0.093	0.091	0.085	0.109	0.129

Notes: This table reports the regressions of the quality-related measurement error against various measures of initial conditions. Robust standard errors are in parentheses. ***: significant at the 1% level; *: significant at the 5% level; *: significant at the 10% level. Each observation is a prefecture pair. The dependent variable is $|T_{ij,\text{uniform}}^m/T_{ij,\text{baseline}}^m - 1|$, where $T_{ij,(\cdot)}^m$ is the travel time between prefectures i and j in 2017 via mode m under the quality-adjusted “baseline” and the unadjusted “uniform” measures of infrastructure. “One Remote” indicates either i or j (but not both) has an average initial travel time above the median, and “Both Remote” indicates that both i and j ’s initial travel times are above the median. The reference group is “None Remote”: both i and j have below-median initial travel time. “Ln(ini. Pop)” and “Ln(ini. GDP)” are the logarithm of the average initial population and GDP of i and j , respectively.

in real wage, output, and population in the quantitative analysis.

The direction of the correlation between the measurement errors and the initial conditions depends on the exact assumptions of the uniform quality. For example, as reported in Table A.5 in the appendix, under the assumption that all the roads are constructed with the highest possible quality, the biases are larger for the remote and poorer locations, similar to those reported in Table 3; however, under the assumption that all the roads are constructed with the lowest possible quality, the biases are larger for the more central and wealthier locations. The result is not a surprise, as the quality of actual construction in these central and richer locations tends to be higher, thus further away from the assumed low quality. Nevertheless, regardless of the sign, our main message stands: the significant correlation between the measurement errors and the initial conditions suggests that explicitly accounting for the quality variations is important for empirical work.

External Validation using Timetable Data Lastly, we highlight the importance of quality adjustments by comparing the travel time estimations with and without quality adjustments to the observed timetable data from bus and railroad service providers. As we do not use any timetable data in constructing our dataset, this exercise offers an external validation test to check if quality adjustments matter.

The timetable data cover passenger travel on both the road and the railroad network. We

digitized the timetable compilation published in the [Tong Xin \(1999\)](#), which provides the travel time on long-distance buses originating from all major prefecture cities in China. The timetable data on the railroad network come from the official website of China Railways, www.12306.cn. The railroad timetable contains the arrival and departure times at each stop for a subset of passenger services in 2002 and 2010. From this information, we measure the *changes* in passenger travel time on the railroad system between prefecture cities over time.

Table 4: External Validation with the Timetable Data

	Bus Timetable			Train Timetable		
	(1)	(2)	(3)	(4)	(5)	(6)
Travel Time, Baseline	1.601*** (0.084)		1.688*** (0.292)			
Travel Time, Uniform		1.555*** (0.084)	-0.086 (0.283)			
Δ Travel Time, Baseline				0.215*** (0.023)		0.228*** (0.030)
Δ Travel Time, Uniform					0.255*** (0.061)	-0.071 (0.077)
N	1,823	1,823	1,823	534	534	534
Adj.R2	0.667	0.652	0.667	0.143	0.024	0.143

Notes: The tables report the regression of observed travel time from timetables on the estimated travel time from the database compiled in this paper. Robust standard errors are in parentheses. ***: significant at the 1% level; *: significant at the 5% level; *: significant at the 10% level. The dependent variable in columns (1) to (3) is the observed travel time on long-distance passenger bus services in 1999. The dependent variable in columns (4) to (6) is the change in passenger travel time on railroads between 2002 and 2010. Each observation is a prefecture pair. “Baseline” refers to the travel time computed that accounts for quality variations, and “uniform” refers to the travel time computed assuming no quality variations across time and space.

To study the impact of quality adjustments, we regress the observed travel time between prefectures against the travel time estimations with and without quality adjustments; if our dataset with quality adjustments reflects the improvements in travel speed, we expect the coefficient on the adjusted travel time to be significantly positive when controlling the unadjusted travel time. Table 4 reports the results.

The quality-adjusted travel time estimates can better explain the variations in the observed travel time on both the road and the rail network than the unadjusted estimates. On the road system, both travel time estimates highly correlate with the recorded travel time of the long-distance buses, as shown in columns (1) and (2). However, when we control for both the unadjusted and the adjusted travel time and run a horse race between the two measures in column (3), the estimate without quality adjustment is no longer significant, while the coefficient on the quality-adjusted travel time is barely affected. Results are similar on the railroad system. When we regress the

changes in travel time between 2002 and 2010 against the *changes* in the baseline and the uniform travel time estimates, only the coefficient with quality adjustments is significantly positive. We interpret these results as external validation that our dataset with quality adjustment captures the actual variations in travel speeds on the road and rail network and thus better matches the observed travel time distribution in both cross-sections and over-time.¹⁴

To further study the potential errors from ignoring quality, we turn to a structural model that can be used to analyze the aggregate and distributional impacts of transportation networks in the next Section.

3 The Model

3.1 Individuals and the Migration Decisions

The economy contains a mass $\bar{L} > 0$ of individual workers and $J > 1$ geographically segmented cities, indexed by $j = 1, 2, \dots, J$. An outside world (ROW), denoted as $j = 0$, trades with the J cities inside the country. Individuals can migrate between the $j = 1, \dots, J$ cities subject to frictions but cannot move to or from ROW. Firms can trade between all the $J + 1$ locations subject to variable trade costs. Time, indexed by $t = 0, 1, \dots, \infty$ is discrete and infinite. Individuals living in city j at time t obtain flow utilities according to $u_{jt} = \log(\phi_{jt} \cdot c_{jt})$, where $c_{jt} = \left(\int_0^1 (q_{jt}(\omega))^{\frac{\eta-1}{\eta}} d\omega \right)^{\frac{\eta}{\eta-1}}$ is the consumption of a CES aggregation of intermediate goods indexed by ω defined over the real interval $[0, 1]$ with an elasticity of substitution denoted as η . Individual also enjoys location-specific amenity that depends on an exogenously time in-invariant component $\bar{\phi}_j$ and the population size of the city at period t , L_{jt} : $\phi_{jt} = \bar{\phi}_j \cdot (L_{jt})^\beta$, where β captures the congestion elasticity. Individuals living in location j at period t receive the wage rate, w_{jt} , consume in city j with an ideal price index $P_{jt} = \left(\int_0^1 (p_{jt}(\omega))^{1-\eta} d\omega \right)^{\frac{1}{1-\eta}}$.

Migration At the end of the period, the individuals decide their locations in the next period, subject to migration frictions, which depend on policy barriers and the transportation networks at the time t . We interpret the policy element of the migration frictions as the hukou system in China that acts as an entry barrier into location i , denoted as $\bar{\lambda}_{it}$. The *direct* passenger transportation cost is a function of the passenger travel time we have measured in Section 2. While many factors other

¹⁴Table A.6 in the appendix repeats the exercises above for the uniform high and low-quality measures and finds consistent results.

than travel time affect migration costs, we broadly interpret travel time as a proxy for distance and ease of travel, which is widely shown to affect migration decisions in the literature.¹⁵

We assume the direct cost of moving from j to i using mode $m = 1, 2, \dots, M$ at time t , denoted $d_{ijt}^{m\mathbb{P}}$, to be:

$$d_{ijt}^{m\mathbb{P}} = a^{m\mathbb{P}} \cdot \left(T_{ijt}^{m\mathbb{P}}\right)^{h^{\mathbb{P}}}, \quad (1)$$

where $a^{m\mathbb{P}} > 0$ captures the mode-specific travel costs, and $h^{\mathbb{P}}$ the elasticity of migration costs to travel time. The term $T_{ijt}^{m\mathbb{P}}$ comes from the passenger travel time measured in our dataset between i and j at time t via mode m . As this term determines the migration costs faced by individuals in the model, biases led by ignoring quality variations would manifest themselves through the model via this term.

Denote the matrix of direct costs as $\mathbf{D}_t^{m\mathbb{P}}$ where the (i, j) th element is $d_{ijt}^{m\mathbb{P}}$. We follow [Allen and Arkolakis \(2022\)](#) to model the route choice of migrants. To move from j to i , the migrant decides a mode of transportation and then a *route* r^m of K steps. The route incurs a cumulative cost of $\sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}}$, where $r_k^m (k = 1, 2, \dots, K)$ is the location of the k th-step in the route r^m .¹⁶ Let \mathcal{R}_{ij}^m denote all the possible routes from j to i . We abstract away from multi-modal routes so migrants cannot switch transportation modes along a given route. In the context of inter-regional transportation in China, such an assumption is justified as multi-modal transportation only accounts for a negligible share ([Huang and Mu, 2018](#)).

Given the policy barriers and the state of passenger transportation networks at time t , the individual residing in city j at time t observes a vector of idiosyncratic preference shocks toward each destination and route, $\{\varepsilon_{it, r^m}\}_{i=1, r^m \in \cup_{m=1}^M \mathcal{R}_{ij}^m}^J$, where ε_{it, r^m} is *i.i.d* across location, route, time, and individual. We assume that ε_{it, r^m} follows a Type-I extreme value distribution, with the following CDF $F(\varepsilon) = \exp(-\exp(-\varepsilon - \bar{\gamma}))$, where $\bar{\gamma}$ is the Euler's constant. Lastly, the individual discounts the future with a rate δ . Taking into account all the components described above, the migration

¹⁵See [Tombe and Zhu \(2019\)](#), and [Ma and Tang \(2020\)](#) for the importance of distance in bilateral migration in China.

¹⁶Different from the baseline setup in [Allen and Arkolakis \(2022\)](#) where the cumulative costs of a route take the multiplicative format of $\prod_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}}$, we assume an additive format to be consistent with the Type-I Generalized Extreme Value shocks in the dynamic stochastic choice framework in [Artuc et al. \(2010\)](#) and [Caliendo et al. \(2018\)](#), which will be introduced later in this section. See Appendix D.1 in [Allen and Arkolakis \(2022\)](#) for more details of the additive travel costs.

decision for an individual living in j at time t can be formally written as

$$v_{jt} = \log \left(\phi_{jt} \frac{w_{jt}}{P_{jt}} \right) + \max_{i, r^m \in \cup_{m=1}^M \mathcal{R}_{ij}^m} \left\{ \delta \mathbb{E}[v_{i,t+1}] - \bar{\lambda}_{it} - \sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} + \kappa \cdot \varepsilon_{it, r^m} \right\}, \quad (2)$$

where v_{jt} is the lifetime utility of the individual currently living in location j , and the expectation is taken over the future realizations of the idiosyncratic preference shock.

Compared to the dynamic migration framework in [Caliendo et al. \(2019\)](#), we differ in two aspects. First, [Caliendo et al. \(2019\)](#) assumes time-invariant migration costs, and we allow the migration costs to be time-variant to capture the changes in transportation networks and policy barriers. Second, we embed the route-choice problem from [Allen and Arkolakis \(2022\)](#) into the dynamic migration model, so the individuals in our model are not choosing a destination prefecture but a pair of destination and route at the same time. Despite the changes in the optimization problem, we show that the dynamic migration problem adopts a similar solution compared to [Caliendo et al. \(2019\)](#) and thus, the model can be solved using the standard iterative algorithms.

Let $V_{jt} \equiv \mathbb{E}[v_{jt}]$ denote the expectation of the lifetime utility; we show in [Appendix C.1](#) that V_{jt} takes the following form:

$$V_{jt} = \log \left(\phi_{jt} \frac{w_{jt}}{P_{jt}} \right) + \kappa \log \left(\sum_{i=1}^J \exp(\delta V_{i,t+1} - \lambda_{ijt})^{1/\kappa} \right), \quad (3)$$

where λ_{ijt} is the *expected* travel costs across all possible modes and routes from j , conditional on moving to i :¹⁷

$$\lambda_{ijt} = \bar{\lambda}_{it} - \kappa \log \left[\sum_{m=1}^M \sum_{r^m \in \mathcal{R}_{ij}^m} \exp \left(- \frac{\left(\sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} \right)}{\kappa} \right) \right]. \quad (4)$$

The expected lifetime utility, V_{jt} , contains two parts: the flow utility in the first term and the option value of living in location j at period t as summarized in the second term. As standard in the dynamic discrete choice literature, we can also derive the migration probability from j to i in period t as:

$$\mu_{ijt} = \frac{\exp(\delta V_{i,t+1} - \lambda_{ijt})^{1/\kappa}}{\sum_{i'=1}^J \exp(\delta V_{i',t+1} - \lambda_{i'jt})^{1/\kappa}}. \quad (5)$$

¹⁷[Appendix C.2](#) provides more details on the solution of λ_{ijt} as a function of travel time and parameters.

In the above expression, κ can also be interpreted as the migration barrier elasticity of bilateral migration. Lastly, the population distribution in the next period is as follows:

$$L_{i,t+1} = \sum_{j=1}^J \mu_{ijt} L_{jt}. \quad (6)$$

3.2 Production and Trade

Production The production side of the economy follows [Eaton and Kortum \(2002\)](#): the market structure is perfectly competitive, and every city can produce every variety $\omega \in [0, 1]$. The production function for variety ω in city j at time t takes the form $q_{jt}(\omega) = A_{jt} \cdot \ell_{jt}$, where ℓ_{jt} is the labor input. A_{jt} is the city-specific productivity that depends on the fundamental productivity of the location, \bar{A}_j , as well as a time-varying part that is a function of the population size in the city to capture the agglomeration force with an elasticity of α , $A_{jt} = \bar{A}_j \cdot (L_{jt})^\alpha$.

Internal Trade Costs We model the internal trade costs following [Allen and Arkolakis \(2022\)](#), similar to what we did to passenger travel costs. Conditional on a mode of transportation, m , the *ad valorem* cost of moving directly from j to i at time t is $d_{ijt}^{m\text{ff}} \geq 1$. Further denote the matrix of $d_{ijt}^{m\text{ff}}$ as $\mathbf{D}_t^{m\text{ff}} = [d_{ijt}^{m\text{ff}}]$. We assume that the direct cost of goods transportation is a function of the observed time cost of freight traffic discussed in Section 2, $\mathbf{T}_t^{m\text{ff}}$:

$$d_{ijt}^{m\text{ff}} = \exp \left(a^{m\text{ff}} \cdot \left(T_{ijt}^{m\text{ff}} \right)^{h^{\text{ff}}} \right), \quad (7)$$

where $a^{m\text{ff}}$ captures the mode-specific cost, and h^{ff} the elasticity of trade frictions to transportation time. Similar to the specification of migration costs, quality variations in the measured freight travel time, $T_{ijt}^{m\text{ff}}$, would affect the trade costs in the model and, subsequently, the spatial distribution of economic activity. Different from equation (1), the direct trade costs are *ad valorem*, and therefore require an exponential transformation to ensure $d_{ijt}^{m\text{ff}} \geq 1$. Conditional on mode m , moving from i to j through a route r^m of K steps incurs a multiplicative cost of $\bar{\tau} \cdot \prod_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\text{ff}}$, where $\bar{\tau}$ is the average transportation cost, and r_k^m ($k=1, 2, \dots, K$) is the location of the k th-step in the route r^m . Consumers in i face the following price if they source variety ω from location j through route r^m :

$$p_{ijt, r^m}(\omega) = \frac{w_{jt} \bar{\tau} \cdot \prod_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\text{ff}}}{A_{jt} z_{ijt, r^m}(\omega)},$$

and individuals source from the cheapest location-route. We assume that $z_{ijt,r^m}(\omega)$ is i.i.d. Fréchet shock across time, location, route, and variety and has a shape parameter θ . From these assumptions, it is straightforward to show that the probability that individuals in i source from j along route r^m is:

$$\pi_{ijt,r^m} = \frac{(w_{jt}/A_{jt})^{-\theta} \left(\prod_{k=1}^K \left(d_{r_{k-1}^m, r_k^m, t}^{m\text{ff}} \right)^{-\theta} \right)}{\sum_{n=0}^J (w_{nt}/A_{nt})^{-\theta} \left[\sum_{m=1}^M \sum_{(r^m)' \in \mathcal{R}_{in}^m} \prod_{k=1}^K \left(d_{(r^m)'_{k-1}, (r^m)'_k, t}^{m\text{ff}} \right)^{-\theta} \right]}, \quad (8)$$

and the expected transportation cost between i and j is:¹⁸

$$\tau_{ijt} = \bar{\tau} \cdot \left[\sum_{m=1}^M \sum_{r^m \in \mathcal{R}_{ij}^m} \prod_{k=1}^K \left(d_{r_{k-1}^m, r_k^m, t}^{m\text{ff}} \right)^{-\theta} \right]^{-\frac{1}{\theta}}. \quad (9)$$

Gravity Equation Aggregate equation (8) over m and \mathcal{R}_{ijt}^m leads to the probability that location i sources from j across all possible routes and modes of transportation:

$$\pi_{ijt} = \sum_{m=1}^M \sum_{r^m \in \mathcal{R}_{ijt}^m} \pi_{ijt,r^m} = \frac{(w_{jt}/A_{jt})^{-\theta} \tau_{ijt}^{-\theta}}{\sum_{n=0}^J (w_{nt}/A_{nt})^{-\theta} \tau_{int}^{-\theta}}. \quad (10)$$

The above also implies that the total income in location j is:

$$w_{jt}L_{jt} = \sum_{i=0}^J \frac{(w_{jt}\tau_{ijt})^{-\theta} \left(\bar{A}_j L_{jt}^\alpha \right)^\theta}{\sum_{n=0}^J (w_{nt}\tau_{int})^{-\theta} \left(A_n L_{nt}^\alpha \right)^\theta} w_{it}L_{it}. \quad (11)$$

3.3 The Equilibrium

Define the time-invariant fundamentals as $\bar{\Omega} = \{\bar{A}_j, \bar{\phi}_j\}$, the time-variant fundamentals as $\Omega_t = \{\bar{\lambda}_{jt}, \mathbf{D}_t^{m\mathbb{P}}, \mathbf{D}_t^{m\text{ff}}\}$, and the sequence of endogenous variables as $\Upsilon_t = \{w_{jt}, L_{jt}, V_{jt}, \mu_{ijt}\}$.

Sequential Competitive Equilibrium Conditional on $\{\bar{\Omega}, \Omega_t\}$ and the initial population $\{L_{j0}\}$, a sequential competitive equilibrium of the model is a sequence of Υ_t that solves the following equilibrium conditions:

1. Individuals maximize their life-time utility (3) by choosing a sequence of locations so that equations (5) and (6) hold.

¹⁸See Appendix C.3 for the solution of τ_{ijt} .

2. Firms maximize their profits in each period and trade balance, so that equation (11) holds.

Steady state A steady state of the economy is an equilibrium where both economic fundamentals and endogenous variables are time-invariant.

Congestion We abstract away from congestion issues in passenger and freight transportation emphasized in [Allen and Arkolakis \(2022\)](#) for several reasons. First and foremost, unlike the within-city commuting setup in [Allen and Arkolakis \(2022\)](#) in which individuals travel at a daily frequency and contribute to urban congestion, our passenger transportation refers to low-frequency migration trips that occur a few times a year. These trips rarely contribute to cross-city road congestion and are barely affected by congestion either. Secondly, railroads run on fixed schedules with capacity limits. Unlike road congestion, where travel speed is highly sensitive to traffic volume, travel speed on the rail network is fixed for an extended period and only responds to traffic changes through scheduling changes. Modeling railroad congestion would require factors such as waiting time and endogenous scheduling decisions by rail service providers. While these considerations are essential, it is beyond the scope of the current project.

Nevertheless, omitting road congestion might lead to certain caveats in interpreting our quantitative results. For example, road segments of higher quality tend to be centrally placed, attract larger traffic volume, and might be more congested than remote and low-quality segments. Moreover, the congestion elasticity could also be segment-specific and correlate with the design speed and capacity. We refer the readers to [Alder et al. \(2023\)](#) for more discussions on the issue of highway congestion in China.

4 Quantification

We quantify the model to 291 prefecture-level cities in China, plus one additional location representing the rest of the world (ROW). The 291 cities in our sample are the largest common sample available in which we can observe both the population and economic output dating back to 1995. We set the time at the yearly frequency, and $t = 1$ corresponds to 1995 in the data. We normalize the total population in China to 100, so the population of a location is equivalent to its population share in percentage points. In the rest of this section, we expand the model to include the ROW and then briefly outline the estimation procedure of all the parameters in the model.

4.1 Rest of the World (ROW)

We model ROW as a single location indexed as $j = 0$ with population $L_{0t} = L_0, \forall t$. We set $L_0 = 475$, reflecting the share of China in the world population. As the ROW does not correspond to any single geographical location, we ignore the agglomeration forces and assume the following production function: $q_{0t}(\omega) = \bar{A}_0 \cdot \ell_{0t}$, where \bar{A}_0 is the fundamental productivity of ROW to be calibrated later jointly with the other location parameters.

We model the trade costs between ROW and Chinese prefectures as follows. We first identify the coastal prefectures that directly import and export from the international markets using the transaction-level data from the Chinese Customs and designate these prefectures as the “port cities” (see Table A.7 in the Appendix). We assume an identical and symmetric iceberg trade cost between any port city and the ROW and denote it as τ^* . The iceberg trade costs between a non-port prefecture i and the ROW are then assumed to be $\tau_{i,\text{ROW},t} = \tau_{\text{ROW},i,t} = \tau^* \cdot \tau_{i,j_t(i),t}$, where $j_t(i)$ is the nearest port city to prefecture i in year t , as measured by τ_{ij_t} . The nearest port city could vary over time due to the changes in transportation networks.

4.2 External Parameters

We choose the value for the following parameters from the literature. The agglomeration elasticity, $\alpha = 0.1$ follows Redding and Turner (2015). Both the congestion elasticity $\beta = -0.3$ and the trade elasticity, $\theta = 6.3$ come from Allen and Arkolakis (2022). The migration elasticity, $\kappa = 2.02$, is taken from Caliendo et al. (2019) based on their estimates at the annual frequency. The discount factor, $\delta = 0.97$, reflects an average annual interest rate of 3 percent. The elasticity of substitution in the utility function, η , does not affect any quantitative results after normalization because it only scales the price index and, thus, the welfare level in the equilibrium. We set $\eta = 6$, a value in the middle of the commonly estimated range as reported in Anderson and van Wincoop (2004).

Policy Barrier, $\bar{\lambda}_{jt}$ We map the entry barrier in the migration frictions, $\bar{\lambda}_{jt}$, to the hukou reform index constructed in Fan (2019). In that paper, the hukou reform index summarizes the gradual relaxation of the entry barriers at the prefecture level from 1998 to 2010, based on official announcements. We follow its strategy and extend the index to the 291 prefectures in our sample between 1994 and 2017 using the same data source. Denote the observed hukou index in prefecture j at time t as $k_{jt} \geq 0$, and a higher κ_{jt} means that prefecture j relaxed more policy restrictions and a lower entry barrier. In the end, we model $\bar{\lambda}_{jt} = \bar{\lambda} - \psi \cdot k_{jt}$, in which $\bar{\lambda}$ captures the overall

policy barrier, and ψ the impacts of the hukou reforms. We calibrate these two parameters in the next section.

Table 5: Parameters

(a) External Calibrated				(b) Estimated			
name	value	source	notes	name	value	s.e.	notes
α	0.1	Redding and Turner (2015)	agglomeration elas.	$a^{1,\text{f}}$	0.875	0.003	road costs, freight
β	-0.3	Allen and Arkolakis (2022)	congestion elas.	$a^{2,\text{f}}$	1.052	0.020	rail costs, freight
θ	6.3	Allen and Arkolakis (2022)	trade elas.	h^{f}	0.050	0.010	time elasticity, freight
κ	2.02	Caliendo et al. (2019)	migration elas.	$a^{1,\text{P}}$	0.889	0.028	road costs, passenger
δ	0.97	-	discount factor	$a^{2,\text{P}}$	1.208	0.039	rail costs, passenger
η	6.0	Anderson and van Wincoop (2004)	elas. of substitution	h^{P}	0.200	0.016	time elasticity, passenger

(c) Joint Calibrated			
name	value	target	notes
$\{\bar{A}_j\}$	-	output in 1995	fundamental productivity
$\{\bar{\phi}_j\}$	-	population in 1995	fundamental amenity
$\bar{\tau}$	0.734	internal-trade-to-GDP ratio, 2002	average internal trade costs
$\bar{\lambda}$	16.370	average annual stay rate between 2000 and 2005	average migration costs
τ^*	1.718	average export-to-GDP ratio, 2000 to 2005	international trade barrier
ψ	0.665	average annual stay rate between 2010 and 2015	elasticity of $\bar{\lambda}_{it}$ to the hukou reform index

Notes: This table reports the results of calibration and estimation. The parameters in Panel (a) are calibrated outside of the model. Those in Panel (b) are jointly calibrated on a transition path from 1995 to the long-run steady state, and those in Panel (c) are estimated using the GMM by simulating the same transition path. The asymptotic standard errors are reported in the third column.

4.3 Estimation and Joint Calibration

While the parameters listed in the previous section were pinned down outside of the model, we estimate and calibrate all the other parameters based on the model solutions on the transition path. We adopt a two-layer quantification strategy. On the outer layer, we use the Generalized Method of Moments (GMM) to estimate the parameters that govern the costs and time elasticity of the transportation networks, $\Theta^\chi = \left\{ [a^{m\text{f}}, a^{m\text{P}}]_{m=1}^3, h^{\text{f}}, h^{\text{P}} \right\}$. In the inner layer during the GMM estimation, conditional on a vector of Θ^χ , we jointly calibrate all the other parameters: $\Theta = \{ \bar{A}_j, \bar{\phi}_j, \bar{\tau}, \bar{\lambda}, \tau^*, \psi \}$.

4.3.1 Inner Layer: Joint Calibration

We first describe joint calibration in the inner layer conditional on a vector of Θ^χ from the GMM. At this stage, with a guess of Θ , we have all the information to solve the transition path in levels. In particular, we solve the model for 50 years from the year 1995, which corresponds to $t = 1$. We first compute $\Omega_t = \{ \bar{\lambda}_{jt}, \mathbf{D}_t^{m\text{P}}, \mathbf{D}_t^{m\text{f}} \}$ from the observed data $\{ k_{jt}, \mathbf{T}_t^{m\text{P}}, \mathbf{T}_t^{m\text{f}} \}$ and the parameters

specified in $\{\Theta^\lambda, \Theta\}$ for the years before 2017, at which point $t = 23$. We then assume that the time-varying fundamentals remain at the same levels as in the year 2017 for all the subsequent years. We do not assume that the year 1995 is in an initial steady state. Instead, the initial year is somewhere along a transition path to a future steady state and compute the transition path using a shooting algorithm. With the solution to the transition path, denoted as $\Upsilon_t = \{w_{jt}, L_{jt}, V_{jt}, \mu_{jt}\}$, we then calibrate the parameters in Θ . Appendix C.4 discusses the details of the shooting algorithm.

We first invert the fundamental productivity, $\{\bar{A}_j\}_{j=0}^J$, from the trade balance conditions as specified in equation (11). To do this, we take the w_{jt} and L_{jt} at year 1994, $t = 0$, from the data (see Appendix B.3 for more details), and take the computed $\tau_{ij,0}$ as given. Equation (11) then implies a unique vector of $\{\bar{A}_j\}_{j=0}^J$ (up to a scale) that rationalizes the observed outputs $\{w_{j0}L_{j0}\}$ across locations, including the ROW, in the year 1994.

We recover the fundamental amenity from the data’s population distribution at $t = 1$. Inferring fundamental amenity requires information on future values of $\{V_{jt}\}$ and thus the entire transition path. We solve the model and then match the endogenous vector $\{L_{j1}\}_{j=1}^J$ to the observed population in the year 1995.¹⁹

The average cost of freight transportation, $\bar{\tau}$, governs the volume of internal trade in China. We follow Ma and Tang (2020) and recover $\bar{\tau}$ from an average internal-trade-to-GDP ratio of 0.625 between 2000 and 2005, based on the data from the *Investment Climate Survey* conducted by the World Bank. Similarly, τ^* governs the external-trade-to-GDP ratio. We target an average export-to-GDP ratio of China at 24.5 percent between 2000 and 2005, based on the official statistics of China.

The last two parameters, $\bar{\lambda}$ and ψ , govern the migration costs matrix. The average cost, $\bar{\lambda}$, targets an annual aggregate stay-rate of 98.9 percent between 2000 and 2005.²⁰ The parameter ψ affects the impacts of hukou reform on internal migration. In the data, the annual aggregate stay rate between 2010 and 2015 had declined to 97.7 percent one decade later as computed using the 2015 *One Percent Population Survey*. Such a decline in aggregate stay rate comes from the improvements in the transportation networks and the gradual liberalization of the hukou policies. We use ψ to target the average aggregate stay rates in the later decade between 2010 and 2015.

¹⁹Kleinman et al. (2023) provides a method that backs out fundamental amenity from the observed distribution of population without solving the model in their Online Appendix B.9. To implement their methods, one needs population data from three consecutive years. Such data in China suffer from many quality issues (see Appendix B.3), so we opt to use a computationally heavier method that is more robust to measurement errors in the data.

²⁰This target comes from the 5-year stay-rate of 94.4 percent as computed using the 2005 *One Percent Population Survey*. Note that $0.944^{1/5} \approx 0.989$.

Intuitively, to identify ψ , we assume that all the residual changes in aggregate stay rate between the 2000s and the 2010s, conditional on the changes in transportation networks, come from the policy reforms.

4.3.2 Outer Layer: Generalized Method of Moments

The parameters in the outer layer, Θ^χ , govern the relative costs of different modes of freight and passenger transportation. We estimate these parameters by matching the moments of the mode-specific traffic in the data. From the *City Statistical Yearbooks of China*, we observe the freight and passenger traffic volumes by prefecture-mode each year. Based on these data, we compute two sets of moments: 1) the average traffic share in each year that identifies $\{a^{m^f}, a^{m^p}\}_{m=1}^M$, and 2) the coefficient of variation of total traffic volume in each year that identifies $\{h^f, h^p\}$.

Moments The first set of moments is the average share of traffic going through the road and rail networks annually across prefectures. The relative traffic shares help to identify $\{a^{m^f}, a^{m^p}\}_{m=1}^M$: conditional on $\{\mathbf{T}_t^{m^\chi}\}_{m=1}^M$, a mode with higher a^{m^χ} is more expensive to travel on and therefore see less usage in the data. In freight and passenger transportation, we normalize waterway transportation costs so that road and rail transportation are identified relative to waterway transportation.²¹

We can compute the counterparts of average traffic shares as functions of the transportation networks and the endogenous trade and migration matrix in the model. Denote the share of all the sales from j via mode m in year t , $s_{jt}^{m^f}$, the moment conditions on the freight network is the average shares across years between 1995 and 2017: $\left\{ \sum_{j=1}^J s_{jt}^{m^f} / J \right\}_{t=1}^{23}$. Similarly, denote the share of population flow from j to all the other prefectures via mode m can be computed as $s_{jt}^{m^p}$, the associated moment conditions are defined as $\left\{ \sum_{j=1}^J s_{jt}^{m^p} / J \right\}_{t=1}^{23}$.²²

The second set of moments is the coefficient of variation (CV) of the total freight and passenger traffic across prefectures each year. These moments identify $\{h^f, h^p\}$, the parameters that control the variations in τ_{ijt} and λ_{ijt} conditional on $\{\mathbf{T}_t^{m^\chi}\}$.²³ Intuitively, as h^χ declines, the resulting trade or migration costs matrix becomes uniform and less dependent on the underlying geography

²¹In particular, we set $a^{3,f} = 1$ and $a^{3,p} = 10$. These parameters do not affect any of the quantitative results. The baseline costs for passenger transportation are large to ensure that the spectral radius of the matrix $\mathbf{F}_t^{m^p}$ is smaller than 1.

²²The asymmetry in the migration costs and entry barriers do not affect the mode choice because the mode choice is conditional on a destination.

²³We use the ‘‘coefficient of variation’’ instead of ‘‘standard deviation’’ because, in the data, the total freight and passenger traffic are recorded in physical units of ‘‘tonnes’’ and ‘‘passengers’’, respectively. The mapping of units between the data and their model counterparts is unclear. In this context, the CV is more suitable because it is scale-free.

summarized in $\{\mathbf{T}_t^{m\chi}\}$. Everything else being equal, the resulting variations in trade or migration flows will decline, leading to a smaller CV. Appendix D.1 provides the details on how to compute the moment conditions in the model. Appendix D.2 discusses the details of estimation procedure.

Identification and Results Figure A.4 in the appendix presents the goodness-of-fit graph of the estimation. Overall, the estimation procedure fits the data moments reasonably well, as most data and the model-generated moments cluster around the 45-degree line.

As is common in structural estimation, the parameters are jointly identified by all the moment conditions. Nevertheless, we can still map the identification of each parameter to specific subsets of the moment conditions as suggested by the gradient matrix illustrated in Figure D.2 in the Appendix. In short, the average usage shares of traffic by modes mainly pin down the costs associated with road and railroad transportation, while the variations in traffic volumes across space identify the elasticity parameters, h^f and h^p . We refer the readers to Appendix D.2 for more details.

Panel (b) in Table 5 summarizes the estimation results. Our estimation of the time elasticity of passenger and freight transportation costs is broadly consistent with the literature.²⁴ For both freight and passenger transportation, the road networks are cheaper to use than the rail system, as reflected in $a^{1,\chi} < a^{2,\chi}, \forall \chi$. This is because most freight and passenger traffic goes through the road system. For example, on average, 79 percent goes through the road system for freight transportation. Only 14 percent of the traffic goes through the rail system. Passenger transportation relies more on the road system, with 91 percent of the traffic leaving only 7 percent on the rail.²⁵

²⁴For example, the time elasticity of passenger travel costs in the context of commuting is 0.07 in Ahlfeldt et al. (2015), and 0.14 in Allen and Arkolakis (2022). Our estimate of 0.2 is slightly higher, as our passenger volume data is in a cross-city context. The time elasticity of freight transportation costs is 0.032 in Egger et al. (2023) and 0.125 in Allen and Arkolakis (2022). The distance elasticity in Donaldson (2018) is 0.135. Our estimate of 0.05 sits within the range of estimates.

²⁵The transportation literature estimates that the monetary costs per ton-kilometers on the railroads are lower than those on the road system, based on the posted freight rates, energy usage, and reliability (see, e.g. Beuthe and Bouffieux, 2008; Kreutzberger, 2008). Our results differ due to the estimation procedure: instead of directly computing the costs, we infer the implied costs of a mode from its observed usage. While on a per ton-kilometer basis, the rail systems are cheaper, in reality, it is not widely used due to their limited capacity and coverage relative to highway transportation. The estimations based on freight rates do not factor in the capacity and coverage limits, while our approach interprets these limits as implicit costs of using the rail network. Allen and Arkolakis (2014) also estimates usage costs from observed usage and finds that while the marginal costs of rail usage are lower than that of the road, the fixed costs of railroad travel are significantly higher. As we do not model fixed costs, the low usage of the railroad is again a reflection of a higher marginal cost.

5 Quantitative Results

In this section, we present the quantitative results based on the estimated parameters from the previous section. We start by briefly assessing the welfare and the distributional impacts of the transportation network improvement over the entire period. We then highlight that the model predictions, especially those related to the distributional impacts, are sensitive to the variations in the quality measure of the infrastructure: ignoring quality variations could lead to substantial and persistent bias in the model predictions.

We evaluate the impacts of transportation networks by comparing two sets of simulations, “baseline” and “no-change”. In both simulations, we start with the initial population distribution in 1995 and solve the transition path 50 years forward until 2044. In the “baseline” simulation, the τ_{ijt} and λ_{ijt} matrices are based on the existing transportation networks and policy barriers each year between 1995 and 2017 and fixed to their 2017 levels for all the subsequent years. In the “no-change” counterfactual, we fixed the τ_{ijt} and the λ_{ijt} matrices to their levels in 1995. In other words, the baseline case simulates the economy that converges to a steady state as defined by the transportation networks in 2017. The “no-change” case is the transition path toward a steady state implied by the initial conditions in 1995. Comparing the results reveals the impacts of expanding the transportation networks.

To further decompose the impacts of the transportation networks through trade and migration, we also simulate two other counterfactual transition paths, “ τ -only” and “ λ -only”. As the names suggest, in the “ τ -only” case, we fix the values of λ_{ijt} to the 1995 levels and allow τ_{ijt} to evolve in the same way as in the baseline. In the “ λ -only” case, we do the reverse and only allow λ_{ijt} to change over time.

5.1 Aggregate and Distributional Impacts

Aggregate Impacts The expansion of transportation networks significantly improved aggregate welfare in the long run by about 57 percent.²⁶ Figure 6 summarizes the impacts of transportation networks on aggregate welfare. In this figure, we plot the transition paths of aggregate welfare in the four counterfactual simulations and normalize the initial levels in 1995 to 1. In the “no-change” case, the aggregate welfare grows by 4.7 percent—a benchmark against which we compare all other counterfactual cases. The aggregate welfare in the “no-change” simulation grows because

²⁶We define welfare of people living in location i at time t as the amenity-adjusted real wage: $\frac{w_{it}}{P_{it}} (L_{it})^\beta \bar{\phi}_i$. The aggregate welfare at time t is the average welfare across locations weighted by L_{it} .

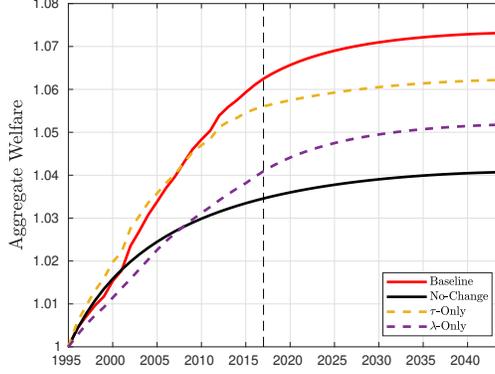


Figure 6: The Aggregate Impacts of Transportation Networks

Notes: This figure plots the transition paths of national welfare in China in four simulations. “Baseline” refers to the case in which both τ and λ matrices evolve based on the actual data between 1995 and 2017, then fixed to their 2017 levels in subsequent years, respectively. “No-Change” is the case in which both τ and λ are fixed to their 1995 levels. “ τ -only” is the case in which only λ is fixed at their 1995 level, and “ λ -only” is the case in which only τ is fixed at their 1995 level. In all the cases, the output level is normalized to 1 in the initial year 1995. The vertical dashed line indicates 2017—the last year we have available data on transportation networks.

the initial year is transitioning to a long-run steady state. Along this transition path, workers migrate to prefectures with higher productivity and amenities, leading to higher aggregate welfare. In the baseline case, the aggregate welfare grows by 7.4 percent, which is $7.4/4.7 - 1 \approx 57$ percent higher.²⁷ The aggregate return to transportation comes from both trade and migration. From the trade channel, better infrastructure allows for higher gains from trade. At the same time, lower migration frictions facilitate population flows toward more productive locations.

The impacts on aggregate welfare imply an annualized rate of return to the investments in transportation networks between 3.0% and 6.8%. As explained in detail in Appendix D.3, we estimate that in the sample period, China invests, on average, 5.6% of GDP each year on transportation networks. Using the gap in GDP growth rates between the baseline and the counterfactual economy without transportation network expansion, we estimate that by 2017, the investments had already yielded an accumulated return of 92.2% or an annualized return of 3.0%. Extending the scope of accounting to 2044 leads to an accumulated return of 2242% over 50 years or 6.8% per year. The finding is similar to those in Wu et al. (2021), who estimated the return on infrastructure investment using firm-level data to be 6 percent.

The long-run impacts of transportation networks mainly come from trade liberalization. The

²⁷The model is not designed to capture economic growth in the data as we have abstracted away from many factors that lead to higher aggregate total factor productivity (TFP), such as technology improvements, physical capital formation, and human capital investments. As a result, the changes in aggregate welfare are modest compared to the data. Nevertheless, we show that transportation networks could lead to a substantial difference in the growth rates of aggregate welfare.

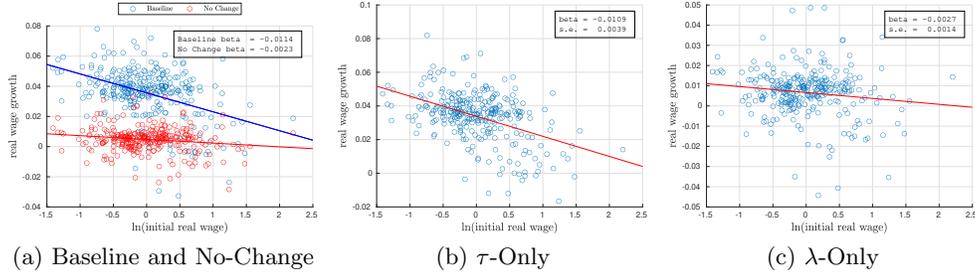


Figure 7: Transportation Networks On β -Convergence

Notes: The figures present the impacts of transportation network improvements on β -convergence between 1995 and 2017 of the real wage. The β -convergence is the coefficient of regressing the growth rate of a variable between 1995 and 2017 against the logarithm of its initial levels in 1995. Each dot represents a prefecture.

two dashed lines in Figure 6 report the transition paths of the τ -only and the λ -only cases. Towards the end of the simulations, the reductions in τ alone lead to a 6.9 percent welfare gain, or $(6.9 - 4.7)/(7.4 - 4.7) \approx 80$ percent of the overall welfare gain led by transportation networks expansion; the λ -only case leads to a 5.3 percent welfare gain, which contributes to about $(5.3 - 4.7)/(7.4 - 4.7) \approx 22$ percent of the overall welfare gain. The net impacts of trade and migration add up to $80 + 22 = 102$ percent, which suggests a mild offsetting effect between the two channels.²⁸

Distributional Impacts We measure spatial inequality using real wage convergence, commonly known as “ β -convergence” in the macroeconomics literature following Barro and Sala-i-Martin (1992). In particular, we regress the real wage growth rate between 1995 and 2044 against the initial level in 1995 in each prefecture. A negative coefficient suggests that the richer prefectures in 1995 grew slower than the poorer ones, so spatial inequality is expected to decline over time. Figure 7 summarizes our findings. Appendix Table A.9 reports consistent results using alternative measures of inequality.

The expansion of transportation networks has significantly reduced spatial inequality. In the baseline case, as shown in the first panel of Figure 7, the distribution of economic activity measured by real wage exhibits strong convergence as the β parameters are significantly negative at -0.0115 . In stark comparison, without the changes in the transportation networks, spatial inequality would have stayed roughly the same as in 1995: the convergence parameter in the “no-change” case is only -0.0027 . In other words, around $1 - 0.0027/0.0115 \approx 77$ percent of the real wage convergence

²⁸The short-run return to migration liberalization could be negative, as shown in Figure 6. In the long run, the return to the reduction in λ starts to turn positive. The delayed return to migration liberalization comes from the forward-looking behavior of the migrants, as discussed in detail in Appendix D.4.

in China comes from the changes in transportation networks.²⁹

The reduction in internal trade frictions is the primary driver of regional convergence. To study the source of regional convergence, we again decompose the impacts into the “ τ -only” and the “ λ -only” cases and report the results in the second row of Figure 7. If better infrastructure only facilitates inter-city trade, the convergence parameter of real wage stands at -0.0113 , which is around $0.0113/0.0115 \approx 98$ percent of the convergence in the baseline case. In comparison, in the λ -only case, the convergence parameter is only -0.0028 , similar to the “no change” case. The role of internal trade as an equalizer is not a surprising finding, as better market access often disproportionately benefits poorer and remote locations, as explained before. Figure A.5 in the Appendix reports the β convergence plots for output and population, respectively. The qualitative message remains the same.

5.2 The Impacts of Quality

In this part, we evaluate the importance of measuring the quality of the infrastructure over time and space. To do so, we carry out another set of counterfactual simulations using the travel time estimates under the assumption of uniform quality, the same as those used in Figure 5. Comparing the model predictions between this “uniform” counterfactual and the baseline results highlights the impacts of quality variations. We summarize the results in Table 6 and Figure 8.

Panel (a) in Table 6 summarizes the relative errors in the real wage, output, and population growth rates at the prefecture-level due to ignoring quality variations. To be precise, we define the relative error for prefecture i at time t for variable x as $|g_{it}^{\text{uniform}}(x)/g_{it}^{\text{baseline}}(x) - 1|$, where $g_{it}^{\text{baseline}}(x) \equiv x_{it}/x_{i,1}$ is the growth rate of variable x between time t and the initial period, using the baseline estimates of travel time that takes into account quality differences. Similarly, $g_{it}^{\text{uniform}}(x)$ is the prefecture-level growth rates computed using the travel time estimates assuming uniform infrastructure quality across space and time.

Ignoring quality variations leads to substantial biases in key model predictions, as shown in the top panel of Table 6. For example, the median error of the real wage growth rates is 33% at $T = 10$ and persists in the steady state at 19%. The relative errors increase considerably at the right tail of the error distribution. At the 90th percentile, the relative error ranges between 33% to

²⁹Our results contrast with findings in Kleinman et al. (2023) for the case of the U.S. In their paper, the authors showed that the initial conditions in the U.S. during the 1960s predicted regional convergence in the next several decades without any future changes in the location fundamentals and transportation networks. The contrast with our results shed light on the country-specific driving force of regional convergence.

Table 6: Measurement Errors due to Ignoring Quality

(a) Relative Errors in Prefecture-Level Growth Rates

	Real Wage Growth Rates				Output Growth Rates				Population Growth Rates			
	50th	90th	99th	Max	50th	90th	99th	Max	50th	90th	99th	Max
T=10	0.33	0.67	5.42	70.30	0.07	0.34	2.68	4.12	0.03	0.16	0.76	3.98
T=20	0.21	0.35	3.56	10.69	0.05	0.28	3.59	31.14	0.03	0.17	2.93	15.03
Steady State	0.19	0.33	1.75	11.08	0.03	0.31	3.36	13.81	0.03	0.20	1.20	5.37

(b) Relative Errors in Spatial Inequality

	$100 \times \beta$ -coef.			$100 \times \Delta$ Gini coef.			$100 \times \Delta$ std(log)		
	Baseline	Uniform	Error	Baseline	Uniform	Error	Baseline	Uniform	Error
T=10	-0.43	-0.60	0.39	-0.18	-0.23	0.29	-0.26	-0.35	0.34
T=20	-0.96	-1.24	0.29	-0.38	-0.46	0.21	-0.57	-0.71	0.24
Steady State	-1.14	-1.43	0.25	-0.43	-0.51	0.19	-0.67	-0.81	0.21

Notes: This table presents the relative errors in prefecture-level growth rates and spatial inequality that come from ignoring the quality of infrastructure. Panel (a) reports the relative errors in growth rates of variable x in prefecture i at time t . The relative error is defined as $|g_{it}^{\text{uniform}}(x)/g_{it}^{\text{baseline}}(x) - 1|$, where $g_{it}^{\text{baseline}}(x) \equiv x_{it}/x_{i,1}$ uses the baseline estimates of travel time, and g_{it}^{uniform} is computed using the travel time estimates assuming uniform infrastructure quality. Column headers refer to the percentiles of the relative error distribution across prefectures. Panel (b) reports the relative errors in spatial inequality of real wage across prefectures. The “Error” column reports the relative error between the “uniform” and the “baseline” simulations. The scaling by 100 is for expositional ease and does not affect the relative error measures.

67%. The maximum error could be as significant as 7030%. Even if the researcher were to exclude the outliers in the error distribution by dropping the top 1% of the error distribution, the relative errors still range between 175% and 542% at the 99th percentile. The median relative errors in predicting output and population growth rates are smaller at 3% to 7%. However, the right tails of the error distributions exhibit similar magnitudes as the real wage. For example, at the 99th percentile, the relative errors range between 268% and 359% for output and between 76% and 293% for population growth rates. Table A.8 in the Appendix shows that the substantial biases in prefecture-level growth rates can also be observed under the alternative measures of errors such as absolute errors or the absolute-errors-to-standard-deviation ratios.

More strikingly, the measurement errors that come from ignoring quality are not random; instead, they correlate with the initial conditions of the prefecture. Figure 8 plots the measurement errors of real wage, output, and population growth rates at the steady state against the initial values of these variables, respectively. Among all three variables, the measurement errors are more pronounced in prefectures with lower initial conditions. In other words, ignoring quality variations tends to overestimate the growth rates of those initially small and less developed locations. The reason for this negative correlation is straightforward: the initially poor locations are more likely to sit on rugged terrain, and thus, their surrounding transportation network tends to have lower

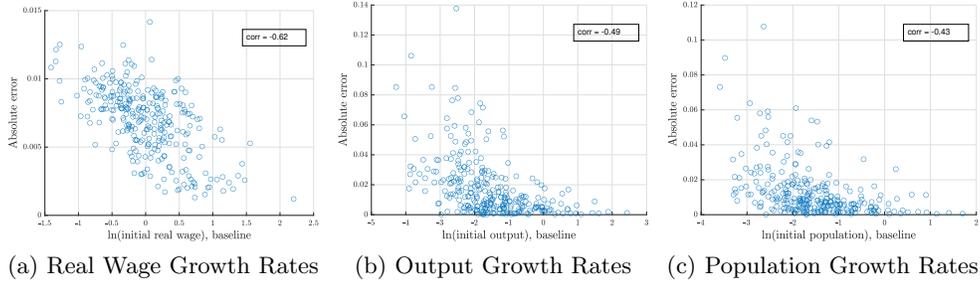


Figure 8: Measurement Error and Initial Conditions, Steady State

Notes: The figures present the relationship between the measurement errors resulting from ignoring quality at the prefecture level against the initial conditions. The vertical axes in these graphs are the “absolute errors” in growth rates due to ignoring quality, which is defined as $|g_{it}^{\text{uniform}}(x) - g_{it}^{\text{baseline}}(x)|$, where $g_{it}^{(\cdot)}(x)$ is the growth rate of variable x between the long-run steady state and the initial state.

design speeds than those in the eastern flood plains.³⁰ Assuming uniform quality overstates the quality of the infrastructure closer to the poor locations and thus overestimates the benefits of such infrastructure. Figure A.6 in the Appendix shows that in addition to the steady state growth rates shown in Figure 8, the negative correlation can also be observed along the transition paths.

The implications of the non-randomness of measurement errors are two-fold. First, empirical researchers shall not treat these errors as classical measurement errors, and therefore one shall not expect the typical solutions, such as instrumental variables, to address the measurement errors led by quality variations. Second, the correlation also implies that the distributional impacts of transportation networks are particularly sensitive to variations in quality. Panel (b) in Table 6 highlights this point by comparing the changes in spatial inequality in real wage between the baseline model and the “uniform” counterfactual. In both exercises, better connectivity always lowers spatial inequality. However, the magnitude differs considerably. For example, the β -convergence coefficient at the $T = 10$ is -0.0043 in the baseline model. The β -coefficient in the “uniform” counterfactual is $0.0060/0.0043 - 1 = 39\%$ higher at -0.0060 . The upward bias slightly alleviates but is still substantial in the long-run steady state at 25% . Similar patterns also emerge under other measures of spatial inequality, as shown in the same table. Ignoring quality differences over-states the distributional impacts of infrastructure improvements because it over-states the beneficial impacts in initially poor locations, as shown in Figure 8. Table A.9 in the Appendix

³⁰The direction of bias in growth rates is consistent with those bias in travel time on the road network reported in Table 3. The similarity between the bias in growth rates and the travel time on the road network comes from the fact that most freight and passenger traffic goes through the road instead of the railroad network. We capture these facts by targeting the share of goods and passenger volume shipped by each transportation mode in the data when estimating the cost parameters, $a^{(\cdot)}$.

reports consistent results for both output and population.

Lastly, ignoring quality variation also leads to a mild positive bias in estimating the aggregate impacts of infrastructure improvements. Table A.10 in the Appendix summarizes the results. In the baseline model, the expansion of transportation infrastructure improves aggregate welfare by 7.3% in the steady state; in comparison, the aggregate welfare impact increases to 7.8% under the “uniform” counterfactual. As the table shows, the bias ranges between 7.4% and 15.8% along the transition path and across different measures of aggregate outcomes. A milder response in aggregate outcomes is expected because the aggregate results are more reliant on the initially rich and populous locations, whose quality of infrastructure does not vary substantially over time.

5.3 Robustness Checks

Table 7 reports the robustness checks of the main results. The first row of the table replicates the main results reported in the previous parts for reference. In the baseline cases, we computed all the measurement errors using the last set of uniform quality measures, which assumes median levels of design speed on mild terrains. As robustness checks, we repeat the exercise using the two alternative uniform quality measures, “uniform high” and “uniform low”. Regardless of the assumptions behind the uniform quality, the measurement errors in growth rates and inequality and the correlation between the measurement errors and the initial conditions remain sizable. The errors are substantially larger under the assumption of the lowest possible quality due to the rapid quality improvements captured in our dataset.

Table 7: Measurement Errors, Robustness Checks

	Real Wage Growth Rates				Output Growth Rates				Population Growth Rates			
	50th	99th	β -err	Corr	50th	99th	β -err	Corr	50th	99th	β -err	Corr
Baseline	0.19	1.75	0.25	-0.62	0.03	3.36	0.08	-0.49	0.03	1.20	0.04	-0.43
Uniform: High	0.12	1.43	0.25	-0.40	0.03	3.90	0.05	-0.29	0.03	1.58	0.06	-0.50
Uniform: Low	0.54	10.22	0.44	-0.60	0.23	31.10	0.44	-0.63	0.18	5.96	0.33	-0.58
w. Capital	0.01	1.02	0.01	-0.40	0.01	0.75	0.02	-0.09	0.02	5.03	0.06	-0.20
$\theta = 8$	0.06	8.03	0.92	-0.35	0.02	1.60	0.10	-0.28	0.02	1.51	0.20	-0.35
$\theta = 4$	0.01	0.31	0.00	-0.06	0.01	0.23	0.00	0.18	0.01	0.23	0.01	-0.03

Notes: This table presents the robustness checks of measurement errors. Columns “50th” and “99th” refer to the percentile in the relative error distribution across prefectures. “ β -err” refers to the error in β -convergence. “Corr” refers to the correlation between the measurement error and the initial conditions. All the variables are computed under the long-run steady state. The first row replicates the baseline results for reference. “Uniform: High (low)” refers to the alternative assumptions on the uniform measure of road quality. θ is the trade and route-switching elasticity. “w. Capital” refers to the version with capital investments.

The baseline model abstracted away from slow-moving factors in each location, such as capital stock. As suggested in Kleinman et al. (2023), location-specific capital accumulation is essential

for analyzing the convergence and the growth rates along the transition path. We present an extended version of the baseline model to include capital accumulation in Appendix E following the framework in Kleinman et al. (2023). We repeat the baseline exercise with and without quality adjustments and report the results in the fourth row of Table 7. The measurement errors are generally smaller than those in the baseline model, although still sizable. For example, instead of a 175% error at the 99th percentile in real wage growth rate, the corresponding error in the model with capital accumulation is “only” 102%. The dampened impact of quality adjustment comes from two reasons: 1) initial capital also affects the real wages, output, and population growth rates, and thus the impacts led by quality variations are less dominant, and 2) as discussed in detail in Appendix E, roughly one-third of the prefectures do not have initial capital stock data and therefore are omitted in the analysis. The prefectures with missing data tend to be remote and poor, thus more susceptible to quality variations. Subsequently, omitting these prefectures thus leads to a smaller bias.

Similar to Allen and Arkolakis (2022), our model features a single parameter, θ , that governs the elasticities of trade, route, and mode choices. In the baseline model, we follow Allen and Arkolakis (2022), interpret the parameter as the trade elasticity, and set it to $\theta = 6.3$. If one were to interpret θ as the route-switching elasticity, the value would be much higher, and if interpreted as the mode-switching elasticity, θ should be smaller than 6.3.³¹ The last two rows of Table 7 report the measurement errors under two alternative values of $\theta = 4$ and 8. Lower θ generally leads to milder measurement errors because the parameter governs the responsiveness of economic variables to trade frictions. As $\theta \rightarrow 0$, variations in transportation infrastructure will be irrelevant for the endogenous variables, and therefore, all relevance of quality adjustment would also be washed out. Nevertheless, even at $\theta = 4$, the measurement errors at the right-tail are still sizable at 31 percent for real wage and 23 percent for output and population growth rates.

6 Conclusion

This paper provides a comprehensive panel dataset that documents the expansion of the transportation networks in China. Our dataset consistently measures the road and railroad quality across time and space compared to the existing measures. Based on this dataset, we evaluate the impacts of transportation networks via a dynamic spatial general equilibrium model that features forward-

³¹See Fan et al. (2021) for the estimation of route-switching and mode-switching elasticity.

looking migration decisions, mode and route choices for freight and passenger transportation, and intercity and international trade. We convey several main messages.

We show that investments in infrastructure have significantly contributed to the economic growth in China, and the returns of trade and migration liberalizations exhibit rich dynamics. We also show that in the case of China, the expansion of transportation networks is the major, if not the sole, driver of regional convergence in the real wage. The equalizing effect of better connectivity mainly comes from the facilitation of internal trade, as it allows better market access for those remote locations.

Moreover, we argue that considering quality variations across time and space is important. Disregarding the quality variations leads to considerable bias in model predictions regarding prefecture-level growth rates. These biases are non-random and correlate with the initial conditions of the prefectures. As a result, ignoring quality differences could lead to a large bias in understanding the distributional impacts of transportation networks.

References

- Ahlfeldt, Gabriel M, Stephen J Redding, Daniel M Sturm, and Nikolaus Wolf, “The economics of density: Evidence from the Berlin Wall,” *Econometrica*, 2015, 83 (6), 2127–2189.
- Alder, Simon, Michael Zheng Song, and Zhitao Zhu, “On (Un)Congested Roads: A Quantitative Analysis of Infrastructure Investment Efficiency using Truck GPS Data,” 2023. working paper.
- Allen, Treb and Costas Arkolakis, “Trade and the Topography of the Spatial Economy,” *The Quarterly Journal of Economics*, 2014, 1085, 1139.
- and —, “The Welfare Effects of Transportation Infrastructure Improvements,” *The Review of Economic Studies*, 02 2022.
- Anderson, James E. and Eric van Wincoop, “Trade Costs,” *Journal of Economic Literature*, September 2004, 42 (3), 691–751.
- Artuc, Erhan, Shubham Chaudhuri, and John McLaren, “Trade Shocks and Labor Adjustment: A Structural Empirical Approach,” *The American Economic Review*, 2010, 100 (3), 1008–1045.
- Asher, Sam and Paul Novosad, “Rural Roads and Local Economic Development,” *American Economic Review*, March 2020, 110 (3), 797–823.

- Banerjee, Abhijit, Esther Dufo, and Nancy Qian**, “On the Road: Access to Transportation Infrastructure and Economic Growth in China,” Working Paper 17897, National Bureau of Economic Research March 2012.
- Barro, Robert J. and Xavier Sala-i-Martin**, “Convergence,” *Journal of Political Economy*, 1992, 100 (2), 223–251.
- Baum-Snow, Nathaniel**, “Did Highways Cause Suburbanization?,” *The Quarterly Journal of Economics*, 05 2007, 122 (2), 775–805.
- , **J. Vernon Henderson, Matthew A. Turner, Qinghua Zhang, and Loren Brandt**, “Does investment in national highways help or hurt hinterland city growth?,” *Journal of Urban Economics*, 2020.
- Berg, Claudia N., Brian Blankespoor, and Harris Selod**, “Roads and Rural Development in Sub-Saharan Africa,” *The Journal of Development Studies*, 2018, 54 (5), 856–874.
- Beuthe, M. and Ch. Bouffloux**, “Analysing Qualitative Attributes of Freight Transport from Stated Orders of Preference Experiment,” *Journal of Transport Economics and Policy*, 2008, 42 (1), 105–128.
- Bonadio, Barthélémy**, “Ports vs. Roads: Infrastructure, Market access and Regional Outcomes,” 2022. NYU Abu Dhabi.
- Caliendo, Lorenzo, Fernando Parro, Esteban Rossi-Hansberg, and Pierre-Daniel Sarte**, “The Impact of Regional and Sectoral Productivity Changes on the U.S. Economy,” *The Review of Economic Studies*, 2018, 85 (4), 2042–2096.
- , **Maximiliano Dvorkin, and Fernando Parro**, “Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock,” *Econometrica*, 2019, pp. 741–835.
- Cameron, A Colin and Pravin K Trivedi**, *Microeconometrics: methods and applications*, Cambridge university press, 2005.
- Donaldson, Dave**, “Railroads of the Raj: Estimating the Impact of Transportation Infrastructure,” *American Economic Review*, April 2018, 108 (4-5), 899–934.
- **and Richard Hornbeck**, “Railroads and American Economic Growth: A “Market Access” Approach,” *The Quarterly Journal of Economics*, 02 2016, 131 (2), 799–858.
- Eaton, Jonathan and Samuel Kortum**, “Technology, Geography, and Trade,” *Econometrica*, Sep. 2002, 70 (5), 1741–1779.

- Egger, Peter H., Gabriel Loumeau, and Nicole Loumeau**, “China’s dazzling transport-infrastructure growth: Measurement and effects,” *Journal of International Economics*, 2023, *142*, 103734.
- Faber, Benjamin**, “Trade Integration, Market Size, and Industrialization: Evidence from China’s National Trunk Highway System,” *The Review of Economic Studies*, 03 2014, *81* (3), 1046–1070.
- Fan, Jingting**, “Internal Geography, Labor Mobility, and the Distributional Impacts of Trade,” *American Economic Journal: Macroeconomics*, 2019, *11* (3), 252–88.
- , **Yi Lu, and Wenlan Luo**, “Valuing Domestic Transport Infrastructure: A View from the Route Choice of Exporters,” *The Review of Economics and Statistics*, 08 2021, pp. 1–46.
- Gurara, Daniel, Vladimir Klyuev, Nkunde Mwase, and Andrea F Presbitero**, “Trends and Challenges in Infrastructure Investment in Developing Countries,” *International Development Policy—Revue internationale de politique de développement*, 2018, *10* (10.1).
- Huang, Sixiang and Dong Mu**, “Discussion on the development strategy of China’s multimodal transport stations,” *Advances in Intelligent Systems Research*, 2018, *161*, 256–262.
- Jedwab, Rémi and Adam Storeygard**, “The Average and Heterogeneous Effects of Transportation Investments: Evidence from Sub-Saharan Africa 1960–2010,” *Journal of the European Economic Association*, 06 2021, *20* (1), 1–38.
- Kapos, Valerie, J Rhind, Mary Edwards, MF Price, C Ravilious et al.**, “Developing a map of the world’s mountain forests.” *Forests in sustainable mountain development: a state of knowledge report for 2000. Task Force on Forests in Sustainable Mountain Development.*, 2000, pp. 4–19.
- Kleinman, Benny, Ernest Liu, and Stephen J. Redding**, “Dynamic Spatial General Equilibrium,” *Econometrica*, 2023, *91* (2), 385–424.
- Kreutzberger, Ekki D.**, “Distance and time in intermodal goods transport networks in Europe: A generic approach,” *Transportation Research Part A: Policy and Practice*, 2008, *42* (7), 973–993.
- Li, Bingjing and Lin Ma**, “JUE insight: Migration, transportation infrastructure, and the spatial transmission of COVID-19 in China,” *Journal of Urban Economics*, 2022, *127*, 103351. JUE Insights: COVID-19 and Cities.
- Liao, Binbing, Jingqing Zhang, Chao Wu, Douglas McIlwraith, Tong Chen, Shengwen Yang, Yike Guo, and Fei Wu**, “Deep sequence learning with auxiliary information for traffic prediction,” in “Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining” 2018, pp. 537–546.

- Lin, Yatang**, “Travel costs and urban specialization patterns: Evidence from China’s high speed railway system,” *Journal of Urban Economics*, 2017, 98, 98 – 123. Urbanization in Developing Countries: Past and Present.
- Ma, Lin and Yang Tang**, “Geography, trade, and internal migration in China,” *Journal of Urban Economics*, 2020, 115, 103181.
- Ma, Liqian**, *The Chronicle of Railway Construction in China (in Chinese)*, China Railway Press, 1983.
- Qin, Yu**, “No county left behind? The distributional impact of high-speed rail upgrades in China,” *Journal of Economic Geography*, 06 2016, 17 (3), 489–520.
- Redding, Stephen J.**, “Goods trade, factor mobility and welfare,” *Journal of International Economics*, 2016, 101, 148–167.
- **and Matthew A. Turner**, “Chapter 20 - Transportation Costs and the Spatial Organization of Economic Activity,” in Gilles Duranton, J. Vernon Henderson, and William C. Strange, eds., *Handbook of Regional and Urban Economics*, Vol. 5 of *Handbook of Regional and Urban Economics*, Elsevier, 2015, pp. 1339 – 1398.
- Rozenberg, Julie and Marianne Fay**, *Beyond the gap: How countries can afford the infrastructure they need while protecting the planet*, The World Bank, 2019.
- Tombe, Trevor and Xiaodong Zhu**, “Trade, Migration, and Productivity: A Quantitative Analysis of China,” *American Economic Review*, May 2019, 109 (5), 1843–72.
- Tong Xin**, *Long-Distance TimeTable Compilation in China*, Tong Xin Press, Beijing, 1999.
- Wong, Woan Foong and Simon Fuchs**, “Multimodal Transport Networks,” 2022. FRB Atlanta Working Paper.
- Wu, Guiying Laura, Qu Feng, and Zhifeng Wang**, “A structural estimation of the return to infrastructure investment in China,” *Journal of Development Economics*, 2021, 152, 102672.
- Xu, Mingzhi**, “Riding on the new silk road: Quantifying the welfare gains from high-speed railways,” 2017. working paper.

Online Appendix

A Additional Tables and Figures

Table A.1: The Physical Maps

Year	National Maps Publisher	Scale	Projection	Provincial Map Publisher
1994	Sino Maps	1:6 million	Albers, 25N, 47E	Sino Maps
1995	N/A			Sino Maps
1996	Sino Maps	1:4.5 million	Albers, 25N, 47E	Global Maps
1997	N/A			Global Maps
1998	N/A			Xi'an Maps
1999	N/A			Xi'an Maps
2000	Sino Maps	1:6 million	Albers	Xi'an Maps
2001	N/A			Xi'an Maps
2002	Sino Maps	1:4.5 million	Albers, 25N, 47E	Xi'an Maps
2003	Sino Maps	1:6 million	Albers, 25N, 47E	Xi'an Maps
2004	N/A			Xi'an Maps
2005	N/A			Xi'an Maps
2006	N/A			Hunan Maps
2007	Guangdong Maps	1:6 million	Lambert, 24N, 46N, 110E	Dizhi
2008	N/A			Dizhi
2009	Sino Maps	1:4.5 million	Albers, 25N, 47E	Renmin Jiaotong
2010	N/A			Dizhi
2011	N/A			Dizhi
2012	Sino Maps	1:4.5 million	Albers, 25N, 47E	Dizhi
2013	Sino Maps	1:4.6 million	Albers, 25N, 47E	Renmin Jiaotong
2014	N/A			Sino Maps
2015	N/A			Sino Maps
2016	N/A			Sino Maps
2017	Sino Maps	1:6 million	Albers	Sino Maps

Non-Map References	
Years	Title
1986-2017	Transportation Yearbooks of China
1994-2017	Railway Yearbooks of China
1881-1981	The Chronicle of Railway Construction in China (Ma, 1983)

Notes: This table presents the basic information on the physical maps in our collection. The Hunan, Xi'an, and Guangdong Maps are regional publishers, while the others are national. The Albers and the Lambert projections are conic projections; the Albers projection is an equal-area projection, while the Lambert projection is conformal. The coordinates following the projections are the reference longitude and latitudes. The coordinate (25N, 47E) is commonly used for Chinese maps as it centers around Henan, the geographical center of China.

Table A.2: The Mapping Between Road Classification and Map Legends

Year	First Rate Road	Highway	Railroad	High Speed Rail
1994	Zhong Yao Gong Lu	Gao Su Gong Lu	Dian Qi Hua Tie Lu	-
1994			Shuang Gui Tie Lu	-
1994			Dan Gui Tie Lu	-
1996	Zhong Yao Gong Lu	Gao Su Gong Lu	Dian Qi Hua Tie Lu	-
1996			Shuang Gui Tie Lu	-
1996			Dan Gui Tie Lu	-
2000	Guo Dao	Gao Su Gong Lu	Tie Lu	-
2002	Guo Dao	Gao Deng Ji Gong Lu	Tie Lu	-
2003	Guo Dao	Gao Deng Ji Gong Lu	Tie Lu	-
2007	Guo Dao	Gao Su Gong Lu	Tie Lu	-
2009	Guo Dao	Gao Deng Ji Gong Lu	Tie Lu	-
2012	Guo Dao	Gao Deng Ji Gong Lu	Tie Lu	-
2013	Guo Dao	Gao Su Gong Lu	Tie Lu	Gao Su Tie Lu
2017	Guo Dao	Gao Deng Ji Gong Lu	Tie Lu	-

Notes: This table shows the correspondence between map legends and the classification of transportation modes used in this paper. We directly report the Pinyin of the legends.

Table A.3: Design Speed (km/h) of Roads and Railroads by Time and Terrain

(a) Road Standards

Revision	Plains	LRH	Hills	Mountains	Plains	LRH	Hills	Mountains
	Highways				First-Rate Roads			
1988	120	120	100	60	100	100	60	60
1997	120	120	120	60	100	100	60	60
2003	120	120	120	80	100	100	80	80
2014	120	120	120	80	100	100	80	80

(b) Railroad Standards, Mixed-Use

Revision	Plains	LRH	Hills	Mountains	Plains	LRH	Hills	Mountains
	National I				National II			
1985	120	100	80	80	100	80	80	80
1999	140	120	80	80	100	80	80	80
2006	160	140	120	120	120	100	80	80
Revision	National III				National IV			
	Plains	LRH	Hills	Mountains	Plains	LRH	Hills	Mountains
1985	80	80	80	80	-	-	-	-
1999	80	80	80	80	-	-	-	-
2012	120	100	80	80	100	80	60	40

(c) Railroad Standards, Freight-Only

Revision	Plains	LRH	Hills	Mountains	Plains	LRH	Hills	Mountains
	Industrial I				Industrial II			
1987	70	70	70	70	55	55	55	55
Revision	Industrial III							
	Plains	LRH	Hills	Mountains	Plains	LRH	Hills	Mountains
1987	40	40	40	40				

Notes: This table summarizes the design speed of the highways and the first-rate roads by revisions of the *Technical Standard of Highway Engineering*, and that of the railroads by various revisions of railroad engineering. The “first-rate roads” is one tier below the highways and are often considered interchangeable with the more commonly known term, “national roads” (Guo Dao). The definition of the terrains is provided in the *Land Regulations in Highway Engineering*, also published by the Ministry of Transportation. “LRH” refers to “Low Rolling Hills”. Please refer to Appendix B.1 for more details.

Table A.4: Uniform Measures of Design Speed

Name	First Rate Road		Speed	Highway		Speed	Railway	
	Speed	Notes		Speed	Notes		Speed	Notes
Uniform High Quality	100	Plains 2014 Rev.	120	Plains 2014 Rev.	160	Plains 2006 Rev.		
Uniform Low Quality	60	Mountains 1988 Rev.	60	Mountains 1988 Rev.	80	Mountains 1985 Rev.		
Uniform Quality	80	LRH, Hills 1997, 2003 Rev.	120	LRH, Hills 1997, 2003 Rev.	100	LRH, Hills 1999 Rev.		

Notes: This table summarizes the definition of design speeds under uniform quality measurements. “Uniform high quality” assumes the highest design speeds, corresponding to those on “plains” in the latest revision. “Uniform low quality” assumes the lowest possible design speed, which is those on “mountains” in the oldest revision. The “uniform quality” assumes the median value across mild terrains (“LRH” and “hill”) from the in-between revisions. The speed for railroads assumes the rate of “National I,” the most common type of mixed-use railroad. The uniform speed in under “uniform quality” is the median of design speeds across terrains and revisions listed under “notes”. Specifically, for first-rate roads, the design speeds on LRH and hills in the 1997 revisions are 100 and 60, respectively, and under the 2003 revisions, 100 and 80. The median of the three numbers (60, 80, 100) is 80. For highways, the design speeds on LRH and hills in both the 1997 and 2003 revisions are 120km/h. For railways, the design speeds on LRH and hills in the 1999 revision are 120 and 80 km/h, and we take 100km/h as the median. The unit for speed in the table is “km/h”, and “LRH” refers to “low-rolling hills”.

Table A.5: Measurement Errors and Initial Conditions, Alternative Measures

(a) Uniform High Quality

	Road		Railroad, Freight		Railroad, Passenger	
	(1)	(2)	(3)	(4)	(5)	(6)
One Remote	0.027*** (0.001)	0.028*** (0.001)	-0.008*** (0.001)	-0.008*** (0.001)	-0.050*** (0.003)	-0.050*** (0.003)
Both Remote	0.041*** (0.001)	0.044*** (0.001)	-0.005*** (0.001)	-0.004*** (0.001)	-0.104*** (0.003)	-0.106*** (0.003)
Ln(ini. Pop)	-0.017*** (0.000)		-0.017*** (0.001)		0.077*** (0.002)	
Ln(ini. GDP)		-0.008*** (0.000)		-0.017*** (0.000)		0.078*** (0.001)
N	39,340	39,340	39,340	39,340	39,340	39,340
Adj.R2	0.138	0.103	0.023	0.029	0.107	0.135

(b) Uniform Low Quality

	Road		Railroad, Freight		Railroad, Passenger	
	(1)	(2)	(3)	(4)	(5)	(6)
One Remote	-0.015*** (0.001)	-0.018*** (0.001)	0.055*** (0.003)	0.051*** (0.003)	-0.056*** (0.008)	-0.059*** (0.008)
Both Remote	-0.046*** (0.002)	-0.053*** (0.002)	0.058*** (0.003)	0.048*** (0.003)	-0.206*** (0.008)	-0.214*** (0.008)
Ln(ini. Pop)	0.043*** (0.001)		0.102*** (0.002)		0.207*** (0.004)	
Ln(ini. GDP)		0.020*** (0.001)		0.081*** (0.001)		0.193*** (0.003)
N	39,340	39,340	39,340	39,340	39,340	39,340
Adj.R2	0.096	0.049	0.095	0.084	0.108	0.124

Notes: This table reports the regressions of the quality-related measurement error against various measures of initial conditions. Robust standard errors are in parentheses. ***: significant at the 1% level; *: significant at the 5% level; *: significant at the 10% level. Each observation is a prefecture pair. The dependent variable is $|T_{ij,\text{uniform}}^m/T_{ij,\text{baseline}}^m - 1|$, where $T_{ij,(.)}^m$ is the travel time between prefectures i and j in 2017 via mode m under the quality-adjusted “baseline” and the unadjusted “uniform” measures of infrastructure. “One Remote” indicates either i or j (but not both) has an average initial travel time above the median, and “Both Remote” indicates that both i and j ’s initial travel times are above the median. The reference group is “None Remote”: both i and j have below-median initial travel time. “Ln(ini. Pop)” and “Ln(ini. GDP)” are the logarithm of the average initial population and GDP of i and j , respectively.

Table A.6: Timetable Data, Alternative Measures

(a) Uniform High Quality

	Bus Timetable			Train Timetable		
	(1)	(2)	(3)	(4)	(5)	(6)
Travel Time, Baseline	1.601*** (0.084)		2.838*** (0.270)			
Travel Time, Uniform		1.803*** (0.103)	-1.459*** (0.337)			
Δ Travel Time, Baseline				0.215*** (0.023)		0.231*** (0.030)
Δ Travel Time, Uniform					0.376*** (0.092)	-0.130 (0.116)
N	1,823	1,823	1,823	534	534	534
Adj.R2	0.667	0.624	0.678	0.143	0.022	0.144

(b) Uniform Low Quality

	Bus Timetable			Train Timetable		
	(1)	(2)	(3)	(4)	(5)	(6)
Travel Time, Baseline	1.601*** (0.084)		1.058*** (0.222)			
Travel Time, Uniform		1.127*** (0.068)	0.395** (0.189)			
Δ Travel Time, Baseline				0.215*** (0.023)		0.226*** (0.030)
Δ Travel Time, Uniform					0.210*** (0.050)	-0.049 (0.062)
N	1,823	1,823	1,823	534	534	534
Adj.R2	0.667	0.657	0.671	0.143	0.025	0.143

Notes: The tables report the regression of observed travel time from timetables on the estimated travel time from the database compiled in this paper. Robust standard errors are in parentheses. ***: significant at the 1% level; *: significant at the 5% level; *: significant at the 10% level. The dependent variable in columns (1) to (3) is the observed travel time on long-distance passenger bus services in 1999. The dependent variable in columns (4) to (6) is the change in passenger travel time on railroads between 2002 and 2010. Each observation is a prefecture pair. “Baseline” refers to the travel time computed that accounts for quality variations, and “uniform” refers to the travel time computed assuming no quality variations across time and space.

Table A.7: List of Port Cities

Tianjin	Dandong	Nantong	Jiaxing	Qingdao	Shenzhen	Jiangmen
Tangshan	Jinzhou	Lianyungang	Fuzhou	Yantai	Zhuhai	Zhanjiang
Qinhuangdao	Shanghai	Ningbo	Xiamen	Weihai	Shantou	Huizhou
Dalian	Suzhou	Wenzhou	Quanzhou	Guangzhou	Foshan	Haikou

Notes: This table lists the 28 prefectures that 1) imports and exports from the international markets in the Chinese Customs database and 2) are on the coast.

Table A.8: Measurement Errors in Prefecture-Level Growth Rates, Alternative Measures

(a) Absolute Errors

	Real Wage Growth Rates				Output Growth Rates				Population Growth Rates			
	50th	90th	99th	Max	50th	90th	99th	Max	50th	90th	99th	Max
T=10	0.01	0.01	0.01	0.02	0.01	0.01	0.04	0.05	0.00	0.01	0.03	0.03
T=20	0.01	0.01	0.01	0.01	0.01	0.03	0.07	0.10	0.00	0.02	0.05	0.08
Steady State	0.01	0.01	0.01	0.01	0.01	0.04	0.09	0.14	0.01	0.03	0.07	0.11

(b) Absolute Errors / STD

	Real Wage Growth Rates				Output Growth Rates				Population Growth Rates			
	50th	90th	99th	Max	50th	90th	99th	Max	50th	90th	99th	Max
T=10	0.78	1.14	1.64	2.75	0.04	0.08	0.21	0.27	0.02	0.05	0.15	0.19
T=20	0.54	0.72	0.93	1.02	0.02	0.08	0.17	0.26	0.01	0.05	0.14	0.21
Steady State	0.43	0.59	0.74	0.85	0.02	0.07	0.14	0.22	0.01	0.05	0.12	0.18

Notes: This table presents the absolute errors in prefecture-level growth rates and spatial inequality that come from ignoring the quality of infrastructure. Panel (a) reports the absolute error, which is defined as $|g_{it}^{\text{uniform}}(x) - g_{it}^{\text{baseline}}(x)|$, where $g_{it}^{\text{baseline}}(x) \equiv x_{it}/x_{i,1}$ uses the baseline estimates of travel time that takes quality differences, and g_{it}^{uniform} is computed similarly using the travel time estimates assuming uniform infrastructure quality across space and time. Note that the absolute errors depend on the norm of the variables in question and therefore are hard to interpret. Panel (b) measures the absolute errors against the standard deviation of $g_{it}^{\text{baseline}}(x)$. For example, the first element in the table implies that the median absolute error is around 72 percent of the standard deviation of real wage growth rates in the baseline. Column headers refer to the percentiles of the relative error distribution across prefectures: the 50th, 90th, 99th, and maximum relative errors.

Table A.9: Measurement Errors in Spatial Inequality, Alternative Outcomes

(a) Output

	$100 \times \beta\text{-coef.}$			$100 \times \Delta$ Gini coef.			$100 \times \Delta$ std(log)		
	Baseline	Uniform	Error	Baseline	Uniform	Error	Baseline	Uniform	Error
T=10	-3.41	-3.77	0.10	-1.44	-1.53	0.07	-1.82	-2.14	0.17
T=20	-6.73	-7.35	0.09	-2.69	-2.89	0.08	-2.88	-3.47	0.20
Steady State	-9.59	-10.32	0.08	-3.54	-3.80	0.07	-2.74	-3.47	0.27

(b) Population

	$100 \times \beta\text{-coef.}$			$100 \times \Delta$ Gini coef.			$100 \times \Delta$ std(log)		
	Baseline	Uniform	Error	Baseline	Uniform	Error	Baseline	Uniform	Error
T=10	-3.42	-3.52	0.03	-2.35	-2.39	0.02	-3.46	-3.66	0.06
T=20	-6.51	-6.83	0.05	-4.37	-4.55	0.04	-5.25	-5.74	0.09
Steady State	-10.09	-10.52	0.04	-5.68	-5.92	0.04	-7.25	-7.93	0.09

Notes: This table reports the relative errors in spatial inequality of output and population across prefectures from ignoring quality variations. We measure spatial inequality in three ways: 1) β -convergence, 2) the changes in the Gini coefficient between time t and the initial period, and 3) the changes in the standard deviation of the natural logarithm between time t and the initial period. The “Error” column reports the relative error between the “uniform” and the “baseline” simulations. The scaling by 100 is for the ease of exposition and does not affect the relative error measures.

Table A.10: Measurement Errors in Aggregate Outcomes

	Aggregate Welfare			Aggregate Output		
	Baseline	Uniform	Error	Baseline	Uniform	Error
T=10	0.031	0.035	0.123	0.021	0.025	0.158
T=20	0.057	0.063	0.091	0.040	0.045	0.108
Steady State	0.073	0.078	0.074	0.051	0.055	0.082

Notes: This table reports the relative errors in aggregate welfare and output growth from ignoring quality variations. The columns with header “baseline” and “uniform” reports the growth rates of aggregate welfare and output between time t and the initial period under the baseline and the “uniform” counterfactual simulations, respectively. The “Error” column reports the relative error between the “uniform” and the “baseline” simulations.

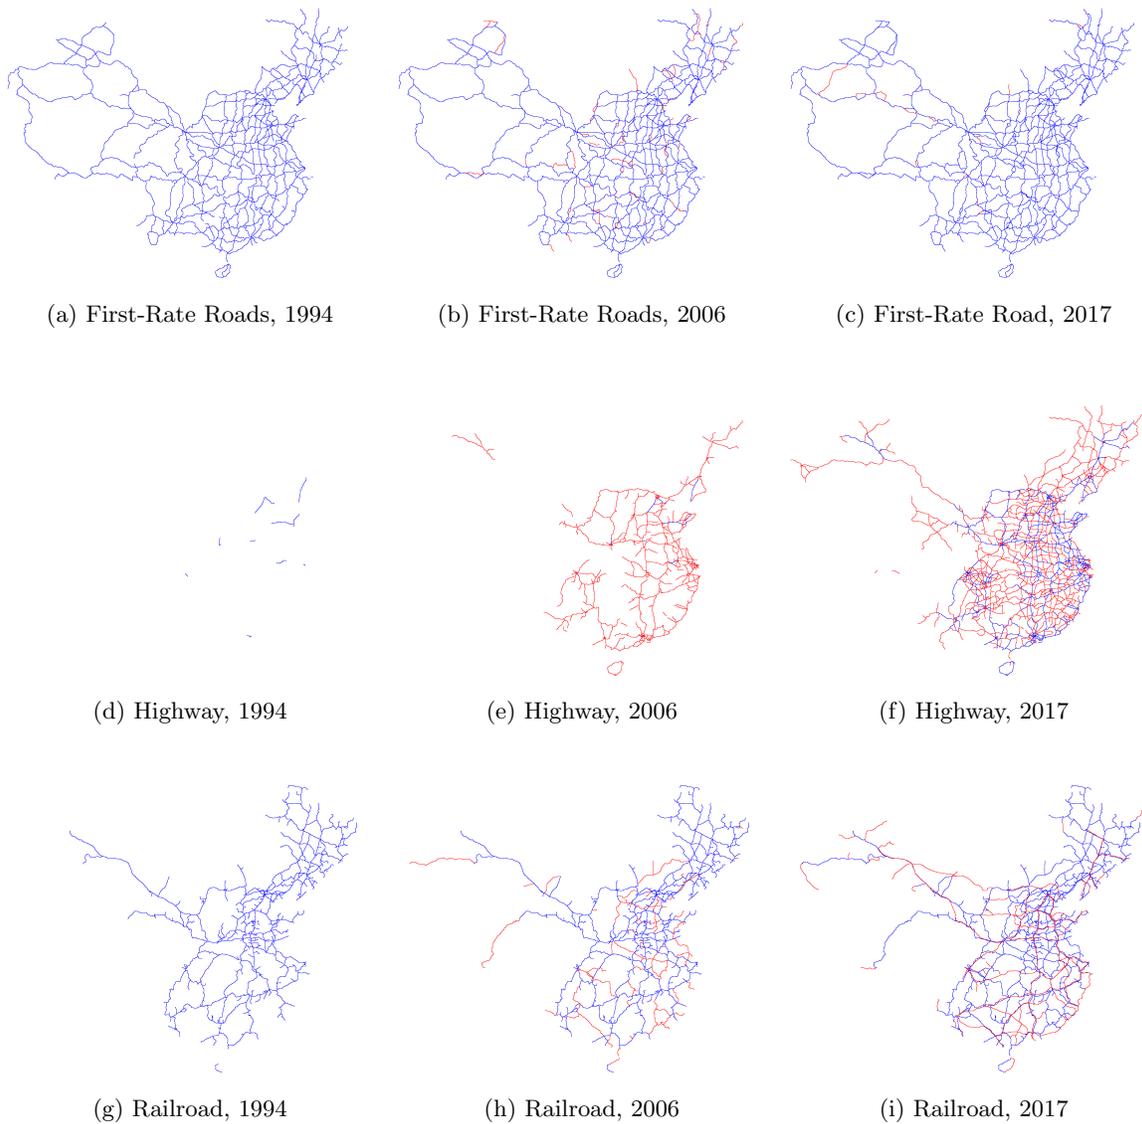
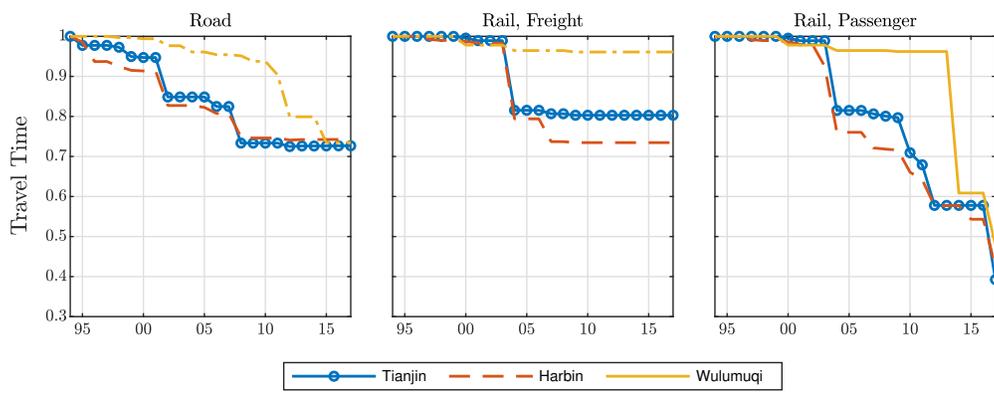
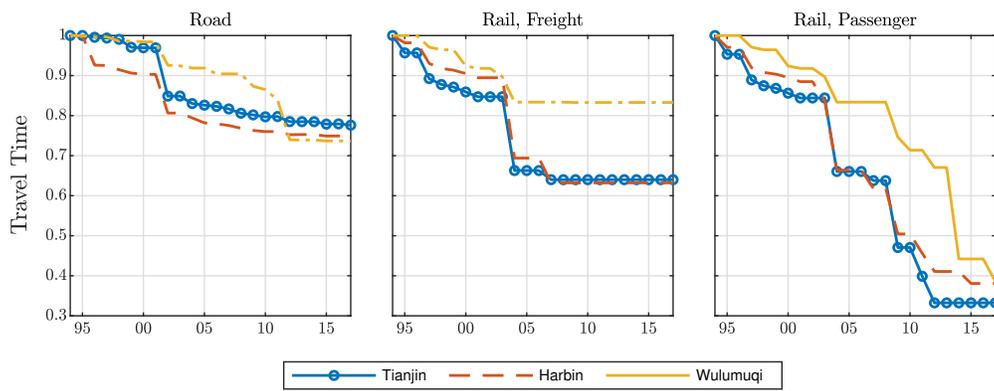


Figure A.1: Changes in Transportation Networks, 1994-2017

Notes: This figure presents the evolution of transportation networks between 1994 and 2017 in several selected years. The “railroad” network includes all the freight-only, passenger-only, and mixed-use railroads. Blue lines indicate existing connections by year t , and the red lines indicate new construction between the year in the previous figure and the current figure (e.g., between 1994 and 2006 in Panel b).



(a) Chengdu



(b) Guangzhou

Figure A.2: The Evolution of Transportation Networks, Distance to Guangzhou and Chengdu

Notes: The figures present the travel time between Chengdu, Guangzhou, and three selected destinations over time. The travel time is normalized to 1 in the initial year. The figure for waterway transportation is omitted because it does not vary over time.

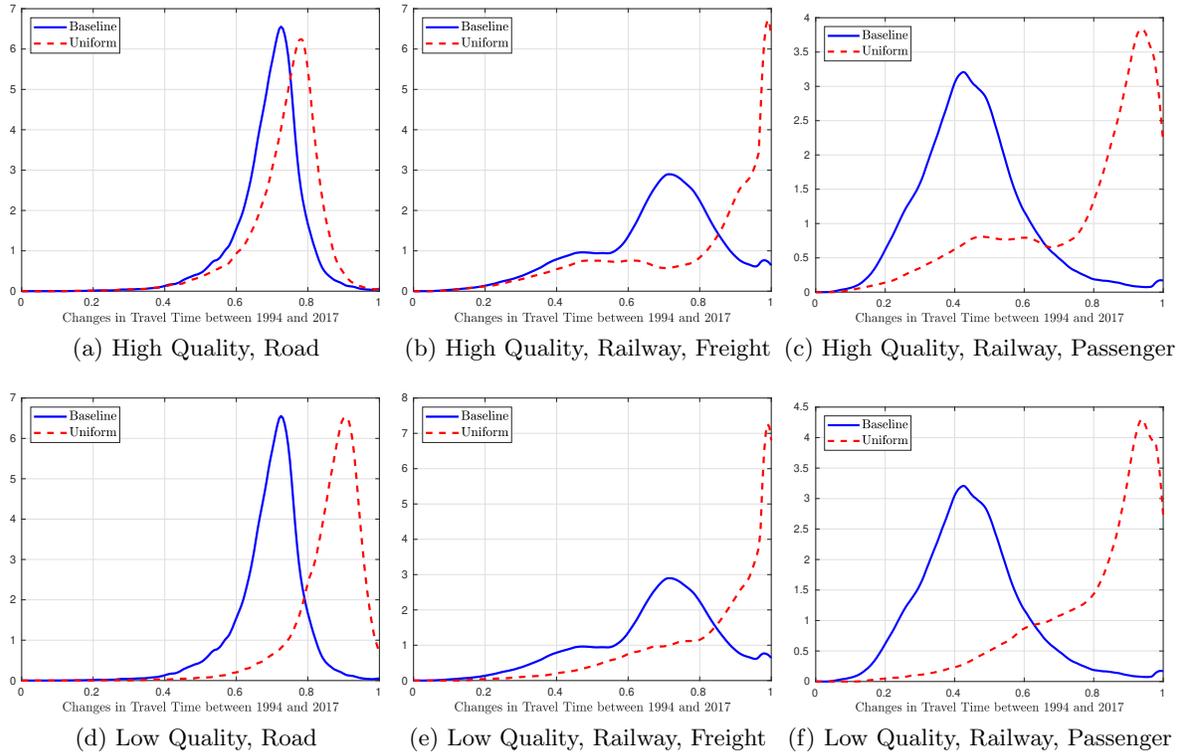


Figure A.3: Comparison between Baseline and Uniform Quality, Alternative Measures

Notes: The figures present the kernel density plot of the ratio of travel time between all prefectures in the year 2017 to those in 1994. The blue line labeled “baseline” accounts for the variations in quality as described in the main text. The red line labeled “uniform” assumes uniform quality across time and space.

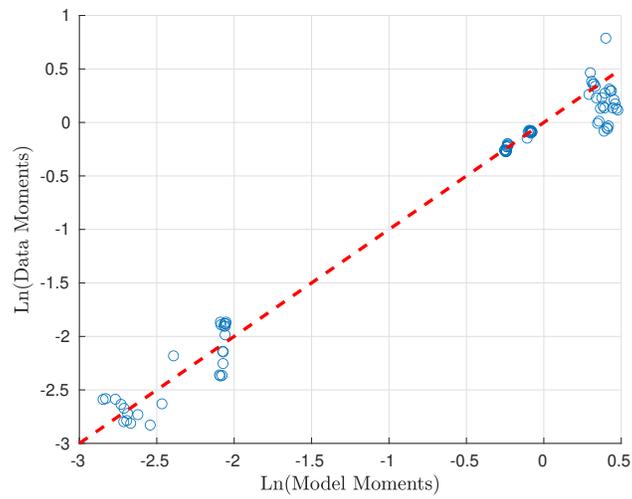


Figure A.4: Goodness-of-Fit, GMM

Notes: This figure plots the goodness-of-fit results of the GMM estimation. The red dashed line is the 45-degree line. The moments in the upper-right corner of the figure are the coefficient-of-variations of freight and passenger traffic each year. Those on the lower-left corner are the shares of freight and passenger traffic via road and rail networks each year.

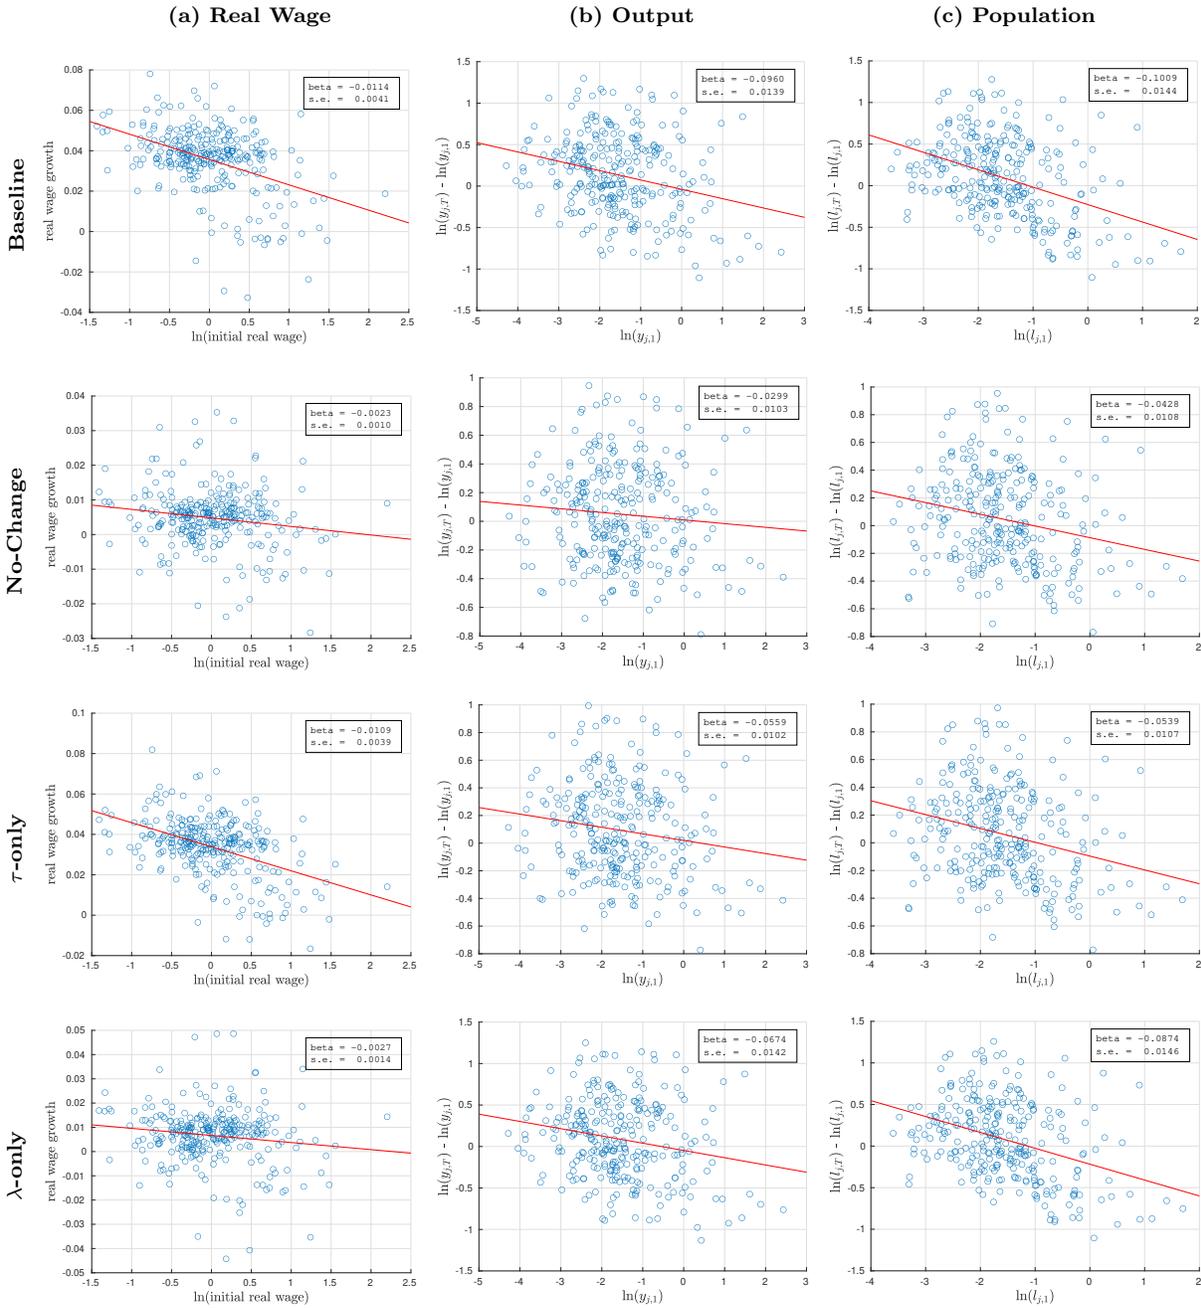


Figure A.5: Transportation Networks On β -Convergence

Notes: The figures present the impacts of transportation network improvements on β -convergence between 1995 and 2017 of real wage, output, and population. The β -convergence is the coefficient of regressing the growth rate of a variable between 1995 and 2017 against the logarithm of its initial levels in 1995. Each dot represents a prefecture.

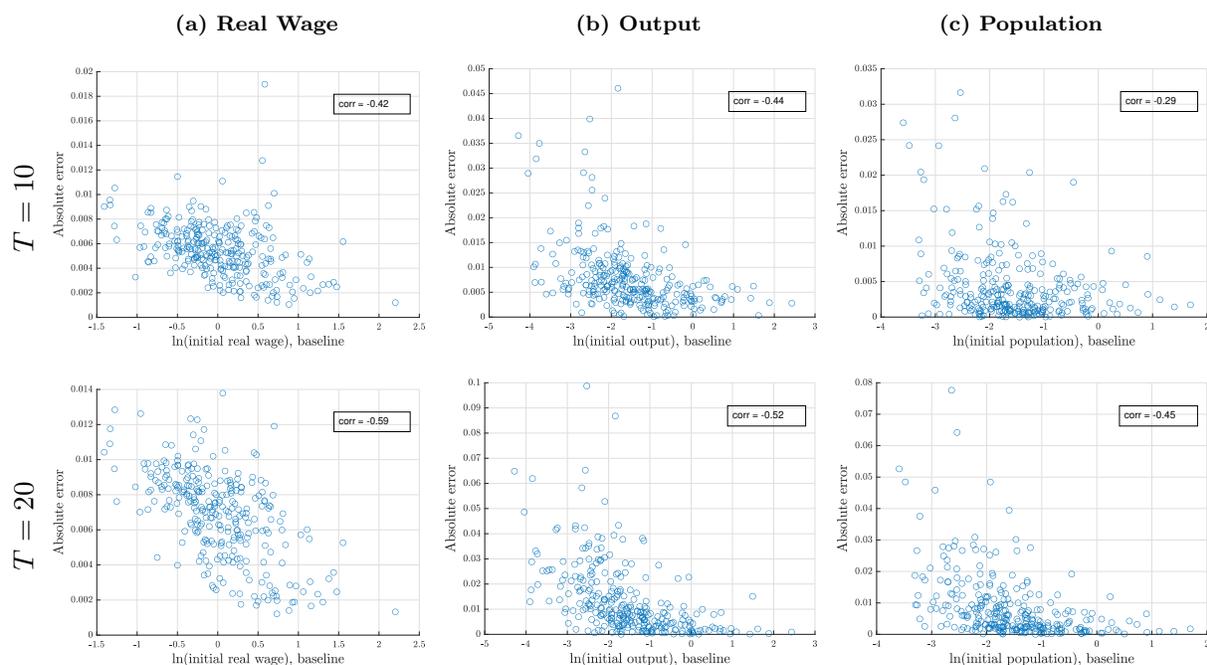


Figure A.6: Measurement Errors and Initial Conditions along the Transition Path

Notes: The figures present the relationship between the measurement errors resulting from ignoring quality at the prefecture level against the initial conditions along the transition path. The vertical axes in these graphs are the “absolute errors” in growth rates due to ignoring quality, which is defined as $|g_{it}^{\text{uniform}}(x) - g_{it}^{\text{baseline}}(x)|$, where $g_{it}^{(\cdot)}(x)$ is the growth rate of variable x between $T = t$ and $T = 1$. In this figure, we set $T = 10$ in the first row and $T = 20$ in the second row.

B Dataset Construction

The temporal coverage of our dataset varies depending on the mode of transportation. For railroads and highways, we track the evolution of every individual railroad and highway from their inception in the country. Specifically, our railroad coverage spans from 1881 to 2017, while our highway coverage extends from 1988 to 2017. However, when it comes to lower-tiered roads known as "first-rated roads," we limit our analysis to the period from 1994 to 2017. This decision is based on the substantial number of lower-tiered roads and the lack of clear documentation regarding their construction in yearbooks and chronicles. For this subset of roads, we assume the existing network in 1994 as a starting point and do not differentiate the pre-1994 roads by vintage.

In the literature, an alternative approach to measuring travel time involves using web-based services such as Google Maps or Baidu Maps. These methods offer the advantage of considering real-time traffic conditions compared to our dataset. However, our dataset has a distinct advantage in terms of historical coverage. Online map services primarily provide travel time estimates based on the current network or, in the case of Baidu Map, an additional snapshot of traffic in 2017 as in [Liao et al. \(2018\)](#). Furthermore, these services do not provide historical transportation network data before 2005, whereas we have traced the development of railroad and highway networks back to their inception.

B.1 Measuring Connectivity

Overview The starting point of compiling the dataset is to collect the published transportation atlas for each year dated back as early as possible. We source the physical maps through several channels, such as libraries, used book dealers, and map collectors. We only choose the national-level atlas with a scale greater than 1:6 million, as smaller maps do not provide enough resolution for color identification. The dimension of the digitized national maps is 12669 pixels in width and 8829 pixels in height, which implies that each pixel is around 500 meters in length in the real world. Out of the 24 years, we have obtained maps for ten years. To fill in the gap years, we resort to the annual *Collection of Provincial Transportation Maps* from various publishers and the transportation yearbooks and chronicles for references. Table [A.1](#) lists the physical maps and other descriptive references in our collection.

One immediate challenge is that the same road is often projected differently across maps, leading to inconsistencies in geo-referencing. The variations in projection come from several sources. First

and foremost, the publishers differ in projection methods. For example, while some publisher, such as *Sino Maps*, uses Albers projection with the reference point at the geographical center of China, other publishers, such as *Guangzhou Publishers*, uses the Lambert projection method with a reference point centered in the province of Guangdong. Moreover, measurement errors also arise due to noise in the maps’ design, printing, and scanning. As we aim to document the evolution of the transportation networks, we must ensure that the existing infrastructure is represented consistently over time.

To ensure consistency over time, we only geo-reference infrastructures based on their first appearance on the national maps. In this way, the same road always has the exact coordinates throughout the panel data. As the national maps only exist in sporadic years throughout the sample period, the year of the first appearance on the national map is not necessarily the year of the construction. To identify the year of the construction, we follow the procedure described below.

In the first step, we deal with the national maps. Denote the ten years in which the national maps exist as the “nodal years” and index them as $t_i, i = 1, 2, \dots, 10$. In our data, the first nodal year is $t_1 = 1994$, and the second nodal year is $t_2 = 1996$. In each national map, we extract four transportation modes by color identification: first-rate road, highway, railway, and water.³² The outcome at this step is the four binary maps by the modes of transportation in each of the ten nodal years.

In the second step, we compare the maps between two consecutive nodal years t_i and t_{i+1} in mode m and construct the annual maps for the years in between. To do so, we treat each binary map as a connected graph and break the graph into “segments”, which are the sets of pixels between the branch points and endpoints of the graph (see Figure B.1 in the Appendix for an illustration). We then compare each segment in year t_{i+1} to that t_i to determine if it existed in t_i or if it is newly constructed between t_i and t_{i+1} . Lastly, for each new construction, we determine the year of the construction by referencing the provincial maps, statistical yearbooks of transportation, railroad, as well as published chronicles on transportation construction in China.³³ We repeat the above process for all the nodal years and all the modes.

At the end of this stage, our dataset identifies the construction year of each geo-referenced segment, from which one can trace the evolution of the connectivity by modes of transportation.

³²Table A.2 in the Appendix provides the correspondence between map legends and the mode of transportation.

³³Across the various reference sources, we first consult the yearbooks and the chronicles because they are official publications with the exact date of the construction. If no yearbooks or chronicles mention the segment, we rely on visual identification from the provincial map collections to determine the year of construction.

Figure A.1 in the Appendix maps the evolution of the connectivity networks in several years. In empirical works, such as in Faber (2014) and Baum-Snow et al. (2020), connectivity is often measured as a binary indicator of whether roads exist on a pixel, or as categorical measures such as highways v.s national roads. Our dataset at this stage is already able to support such analysis. In the next part, however, we show that a binary connectivity measure is insufficient, as the quality of the infrastructure varies significantly over time and space. To address this issue, we discuss the variation and measurement of the quality of the transportation networks in China.

Digitization The digitized physical maps in this project are summarized in Table A.1 and A.2. We use standard procedures to digitize these maps. After scanning the maps, we extract the modes of transportation by color identification and geo-reference each map by inverting the projection methods. The dimension of the digitized national maps is 12669 pixels in width and 8829 pixels in height, which implies that each pixel represents around 500 square meters in the real world.

The transportation networks in the color-identified maps usually have paths 10 to 20 pixels in width. Given that each pixel corresponds to around 500 square meters in the real world, we shrink the color-identified maps down to one-pixel paths using the skeletonization algorithm. We use eight-connectivity at the pixel level, meaning that each pixel in the network is considered connected to all eight of its neighbors, including the diagonal neighbors. We then break the entire skeleton network into “segments”, defined as the set of pixels between the branch and endpoints of the graph. See Figure B.1 for an illustration. The above procedure is applied to all modes of transportation and in all years.

Consistency To ensure that each segment is consistently represented over time, we compare the segments in nodal year t_{i+1} to all the segments in nodal year t_i to determine if the segment already existed in year t_i , or it was constructed between t_i and t_{i+1} . To classify the roads, we first denote segment k in year t_{i+1} as $g_{t_{i+1}}^k$, and then use the following procedure:

1. If there exists a segment in year t_i , denoted as $g_{t_i}^{k'}$ such that $g_{t_{i+1}}^k \subseteq g_{t_i}^{k'}$, then segment k already existed in t_i .
2. If no such segment exists in the previous step, we then manually compare $g_{t_{i+1}}^k$ to the closest segment in t_i^k and determine if $g_{t_{i+1}}^k$ already existed in t_i . To determine the closest segment, we first compute the pair-wise distance between all the pixels in two paths and take the simple

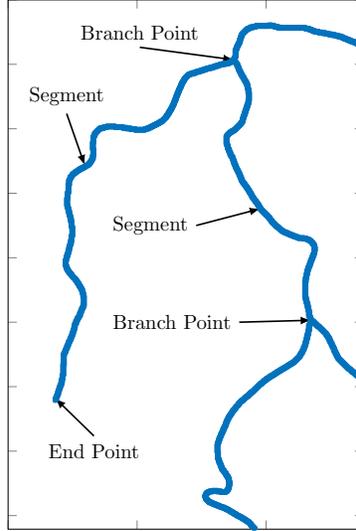


Figure B.1: Definition of Segments, End Points, and Branch Points

Notes: This figure illustrates the definition of “segments” in our dataset. The blue lines indicate a transportation network. A “segment” is the set of pixels between any two branch points or endpoints of the network.

- average. To determine if the segment already existed, we visually compare the segments in the national maps and also refer to the provincial maps, yearbooks, and chronicles.
3. For newly constructed segments, we then refine the segments to avoid overlapped pixels. As shown in the example in Figure B.2 and explained in detail later, a segment in t_{i+1} might contain pixels that are already constructed in t_i . Therefore, we refine the segment $g_{t_{i+1}}^k$ to drop all the pixels that already exists in $\cup_{k'} g_{t_i}^{k'}$.
 4. To determine the year of construction, we first refer to the yearbooks and the chronicles because they are official publications with the exact date of construction and the relevant contexts. The *Transportation Yearbooks* mainly cover the highway and significant railway construction. In contrast, *Railway Yearbooks* and the *Chronicles of Railway Construction* together provide a complete picture of railway construction. If a segment does not show up in any of the yearbooks or chronicles, we then rely on visual identification from the provincial map collections to determine the year of construction. The first-rate roads often fall into this category as they are not deemed important enough to be recorded in the yearbooks.
 5. For railway construction, yearbooks and chronicles could record up to four milestone dates in chronicle order: construction finished (Tong Che), trial-run (Shi Yun Ying), in-business (Yun Ying), and record-closure (Xiao Hao). While the four milestone dates could occur within the

same year, in many cases, they span over two to three years, leading to ambiguities as to what constitutes the “year of construction” of a railway. Ideally, if all four dates are available for every railroad, one would choose a single milestone as the date of construction. Unfortunately, many railways do not have a complete record of all four dates, and therefore we do not have the liberty to designate one type of date as the “date of construction”. As any milestone date could be interpreted as the railway’s completion, we use the first mention of any of the four dates as the date of construction.

Example In this part, we provide an example of identifying a newly constructed highway, as shown in Figure B.2. In this figure, the thin red line is a segment in the second nodal year, $t_2 = 1996$, identified by breaking up the graph in 1996 between all the branch points and endpoints. The thick blue line is the nearest segment in the previous nodal year, $t_1 = 1994$. The figure shows that the two lines overlap, as each is identified separately relative to the network in their respective nodal years. The map shows that the highway in question is the Chengyu Highway, which connects Chengdu and Chongqing in Southwest China. From the yearbooks, the blue line is the early stage of the highway that already existed by 1994. The newly constructed red line is the phase-two construction that was finished in 1995, as recorded in the *Transportation Yearbook of China, 1996*, page 438, translated and quoted here for reference:

“(The province of Sichuan) had made breakthroughs in highway construction. Beating the original plan, the Chengyu Highway started a trial run on July 1st and commenced operation on September 25th (1995). The highway, measuring 340.2 km, constructed in 5 years with a cost of 4 billion RMB, was the very first highway in Sichuan.”

Lastly, we refine the red segment to drop the overlapped pixels with the existing blue segment and assign the construction year of 1995 to the re-defined segment.

B.2 Measuring Quality

Road, Overview Most quality variation in the road network comes from the spatial dimension. In the earlier standards, the design speed of all roads depends on the terrain. For example, in the 1988 revision, the highways are constructed to allow for 120km/h design speed in the plains and low rolling hills (LRH), 100km/h in high hills, and 60 to 80km/h in the mountain regions. Similar

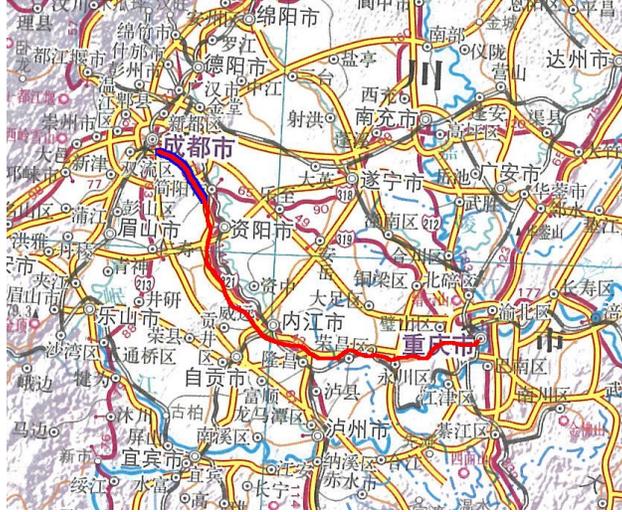


Figure B.2: Identification of New Construction

Notes: This is an example of a newly constructed highway. In the figure, the thin red line is a segment in $t_2 = 1996$, and the thick blue line is the nearest segment in the previous nodal year, $t_1 = 1994$. Based on this and the records in the yearbooks, we determine that the differences between the red and the blue lines form a new segment that was constructed between 1994 and 1996. The highway in question is the Chengyu Highway which connects Chengde and Chongqing. The blue segment is the phase-one construction that was already finished by 1994, and the newly constructed red line is the phase-two construction that was finished in 1995, as recorded in the *Transportation Yearbook of China, 1996*, page 438.

rules apply to lower-rate roads with a lower design speed than highways at all terrains.³⁴

Over the years, the Ministry of Transport has gradually reduced the reliance on terrain factors when designing highways, while still maintaining the terrain dependency for lower-rate roads. For instance, in the 2003 revision, it was mandated that all highways should be designed with a speed limit of 120 km/h, except in particularly challenging regions. Consequently, more elevated highways were constructed in low-rolling hills and hilly terrains.

However, the dependency on terrain factors continued for first-rate roads, commonly known as “national roads.” In plain and LRH regions, the design speed for first-rate roads remained consistently at 100 km/h across all four revisions. In hilly and mountainous areas, there was an increase in the design speed from 60 km/h to 80 km/h in the 2003 revision.

Lastly, the Ministry of Transport (MOT) classifies re-constructed and upgraded roads as new constructions and applies the prevailing standards to these “new” roads, as stated in the Technical Standard of Highway Engineering (1997 Revision). As a result, our dataset also captures the changes in quality that arise from the reconstruction and upgrading of roads. This allows us to

³⁴For example, in 1988, first-rate roads were designed for 100km/h in the plains and 60km/h in the mountains. For the second-rate roads, the design speeds were lowered to 80 and 40km/h, respectively.

track the evolution of road quality over time and incorporate these variations into our analysis.

Road, Details To measure the quality of roads across time and space, we rely on the publications from the Ministry of Transportation from China: the *Technical Standard of Highway Engineering*. We use the following four revisions: 1988 (JTJ01-88), 1997 (JTJ01-97), 2003 (JTG B01-2003), and 2014 (JTG B01-2014).

We apply the design speed from the prevailing design codes at the time to determine the quality of highways and first-rate roads. For example, JTJ01-88 regulates the highways constructed in 1995, and JTJ01-97 applies to those constructed in 1999. As highway construction in China started in 1988, the four revisions covered all the highways constructed in China. The case of first-rated roads is tricky. While we can determine the applicable design codes for the first-rate roads built after 1994, it is impossible to determine the years of construction of those that already existed in 1994 based on our data. When measuring quality, we apply the 1988 revision to all the existing first-rated roads in 1994. This exercise assumes that by assigning a particular road as “the first-rate” in 1994, that road must have a quality closest to the prevailing standards at the moment, the JTJ01-88 revision; lower quality roads would have received a lower rating on the national maps.

The information on design speed comes from the *Technical Standard of Highway Engineering*. In the 1988 revision, the design speed, written as “Ji Suan Xing Che Su Du” as in Pinyin, depends on the terrain as stipulated in Chapter 2.0.2. With engineering difficulty and costs in mind, the technical standard particularly emphasized the dependence on terrain for highways. In the mountainous regions, the default design speed is 80km/h, and the 60km/h speed applies to “challenging segments”. As the technical standard does not provide more detail on the conditions for using the 60km/h design speed, we choose to use 80km/h for all the highways in the mountain areas constructed under the 1988 revision. The design speed for the first-rate roads is precise: 100km/h in plains and low rolling hills and 60km/h in high hills and mountains.

The major change in the 1997 revision is to remove the dependence on terrain for highway construction: “the design speed of highway is no longer linked to the underlying terrain, ..., under normal circumstances, the design speed should be 120km/h”. Under limited conditions, the design speed can be 100 or 80 km/h. In “particularly difficult segments”, 60km/h is still admissible. For this reason, we use a design speed of 120km for all terrains except for the mountains, for which we still use the 80km/h design speed. The construction of first-rate roads is still reliant on the underlying terrain, and there is no change in the design speeds.

The next revision in 2003 focused on first-rate roads. In high hills and mountains, the design speed of this class of roads increased from 60km/h to 80km/h, the same as the highways. In highway construction, the Ministry discourages the 60km/h design, as upgrading a highway with a low-speed design is challenging to accommodate the ever-increasing traffic volume. The 2003 revision still allowed a 60km/h design for highways, but it stipulated that such segments cannot be longer than 15km. The 2003 revision also changed the Chinese translation of “design speed” from “Ji Suan Xing Che Su Du” to a more direct translation of “She Ji Su Du” to be comparable to the international standards.

The design speed of the highways and the first-rate roads did not change in the latest 2014 revision. The latest revision stated that the 2003 revision was already well-crafted and tested in practice. For consistency, the Ministry no longer saw a need to revise the design speeds further.

Railroad, Overview First, the quality variation in the railroad system comes from both the temporal and the spatial dimensions. Similar to the road system, railroads in the hills and mountains limit travel speed. However, different from the road system, the railroads’ design speed changes significantly over time. For example, the design speed of National I railroads on the plains increased from 120km/h in 1985 to 160km/h in 2006. The cross-time improvements in design speed are also seen in the other grades and on the other terrains. The significant changes across time imply that assuming a uniform quality of railroads would underestimate the improvements in connectivity over time.

Second, the operation speed of existing railroads could also change over time, which leads to even higher temporal variation. Unlike roads, where the design speeds are fixed upon construction, the speed of existing railroads changes over time due to continuous engine and route design improvements. During our sampling period, the Ministry of Railroads implemented six waves of speed improvements (“Lu Wang Ti Su”) on the existing rail network, which increased the speed of old networks by as much as 30 percent in some cases.³⁵ We incorporate such significant improvements in the dataset by referencing the yearbooks that provide the details of speed improvements of the major railroads.

Third, different from road transportation, where passengers and goods share the same road, railroad designs vary by intended usage and fall into three categories: passenger-only, freight-only, and mixed-use. The passenger-only railroads, including the High-Speed-Rail (HSR) system

³⁵For example, the design speed of the Longhai Railroad, built in the 1950s, was 120km/h before the first wave and increased to 160km/h after the sixth wave of speed improvements.

introduced in the last two decades, emphasize travel speed but cannot handle heavy loads. On the other hand, freight-only railroads prioritize load capacity over speed. Most of the railroads fall into the last category—the mixed-use railroads. This railroad allows passenger and freight traffic and often balances speed and load. In addition, different from the road system where the legends in the maps identify the rate of the roads (see Table A.2), railroads of all usage and rates are often represented using a single legend in most of the maps. To identify each segment’s intended usage and rate, we again refer to the yearbooks and the chronicles of railroad construction. From these sources, we first recover the railroads’ names to which each segment belongs. We then identify the intended usage and the rate based on their names.

We treat the following two groups of railroads differently. The first group is the HSR system, the high-speed passenger-only railroads. These railroads are designed with speeds much higher than the existing mixed-use networks and are not governed by the engineering standards behind Table 1. The second group is the long-haul freight-only railroads, such as the Dalian-Qinhuangdao railroad, because Panel (c) of Table 1 only covers short-range industrial railroads. Long-haul freight trains typically travel at a faster speed than short-haul industrial railroads. To address these issues, we manually collect the design speed for each HSR and long-haul freight-only railroad from the yearbooks.

Railroad, Details The evolution of railroad engineering design codes is much more convoluted than the highways. Several strands of codes existed, and the classification of rates also changed over time. Table B.1 summarizes the mapping between railroad codes and the rates railroads that they cover. In the rest of this part, we discuss several specific issues related to railroad design codes.

First, unlike highway engineering, the design codes for railway engineering do not map speed to terrain in every revision. The 1999 revision (GB50090-99) was the only one that explicitly mapped design speed and terrain in Chapter 1.0.5, Table 1.0.5-2 for National I and II railroads. Other revisions still emphasize the design speed’s dependency on terrain but do not explicitly provide a mapping between the two. For example, in the 2006 revision, Chapter 1.0.5 defines the speed of National I railroads at 160km/h, 140km/h, and 120km/h, depending on terrain, but does not provide any further indication.

We use the mapping defined in the 1999 revision as a reference and infer the mapping between speed and terrain in the other revisions to address this issue. In principle, whenever the design code

Table B.1: The Mapping Between Railroad Rates and the Codes of Railroad Design

Codes: Revision:	<i>Code for Design of Railway Line</i>				CDRL (III,IV) ^a	CDSRLIF ^b
	1985	1999	2006	2017	2012	1987
Doc.Number:	GBJ90-85	GB50090-99	GB50090-2006	TB10098-2017	GB50012-2012	GBJ12-87
National I:	■	■	■	■	□	□
National II:	■	■	■	■	□	□
National III:	■	■	□	□	■	□
National IV:	□	□	□	□	■	□
Local I:	□	□	□	□	■	□
Local II:	□	□	□	□	■	□
Local III:	□	□	□	□	■	□
Industrial I:	□	□	□	□	□	■
Industrial II:	□	□	□	□	□	■
Industrial III:	□	□	□	□	□	■

Notes: This table lists the mapping between the railway design codes with various rates of railroads. ■ indicates that the design code covers the railroad rate in question, and □ indicates otherwise.

^a Code for Design of III and IV Rated Railway Line.

^b Code for Design of Standard Railway Line for Industrial Firms.

allows for several categories of design speed within a rate, we assign the highest design speed to the plains, followed by LRH, hills, and mountains. If the design code allows for four-speed categories, such as National IV in the 2012 revision, we map the four speeds to the four types of terrain. In the case of three-speed categories, such as National III in the 2012 revision, we assign the same speed to “hills” and “mountains”, following the mapping of National I in the 1999 revision. In the case of two-speed categories, we assign the same speed to LRH, hills, and mountains, following the mapping of National II in the 1999 revision. For example, the 2006 revision stipulated three categories of speed for National I railroads at 160, 140, and 120 km/h. Based on this, we assume the design speed on plains to be 160km/h, LRH to be 140km/h, and hills and mountains to be 120km/h.

Second, different from the road rating system, the rates of railroads had changed significantly over time. In earlier revisions, the Ministry of Railroads rates railroads by “National I, II, and III”, followed by “Local I, II, and III” ratings. The quality of the “Local” rated railroads is typically lower. However, starting from the 2012 revision, local ratings were discontinued, and the local-rated railroads were re-classified as either National III or National IV, a new grade introduced in the 2012 revision. In light of these, we assign the design speed of National III in the 2012 revision to Local I-rated railroads and National IV to Local II and III-rated railroads.

Lastly, similar to the roads built before 1994, we use the rating in 1994 to infer the design speeds of the existing railroads in 1994. The logic is the same as before: if an existing railroad was

designated a particular rating in 1994, its quality must be closest to that rating. Otherwise, the railroad would have received a different rating.

Definition of Terrains As shown in Table 1, the design speed of the roads and railroads differ by four types of terrain: the plains (Ping Yuan), the low rolling hills (LRH, Wei Qiu), the hills (Zhong Qiu), and the mountains (Shan Ling). The four types of terrains are defined in the *Land Regulations in Highway Engineering*, published by the Ministry of Transportation. Panel (a) of Figure B.3 summarizes the definition of terrains used in this paper, and Panel (b) maps the terrains in China as defined. In this part, we provide more details about the definition of terrains.

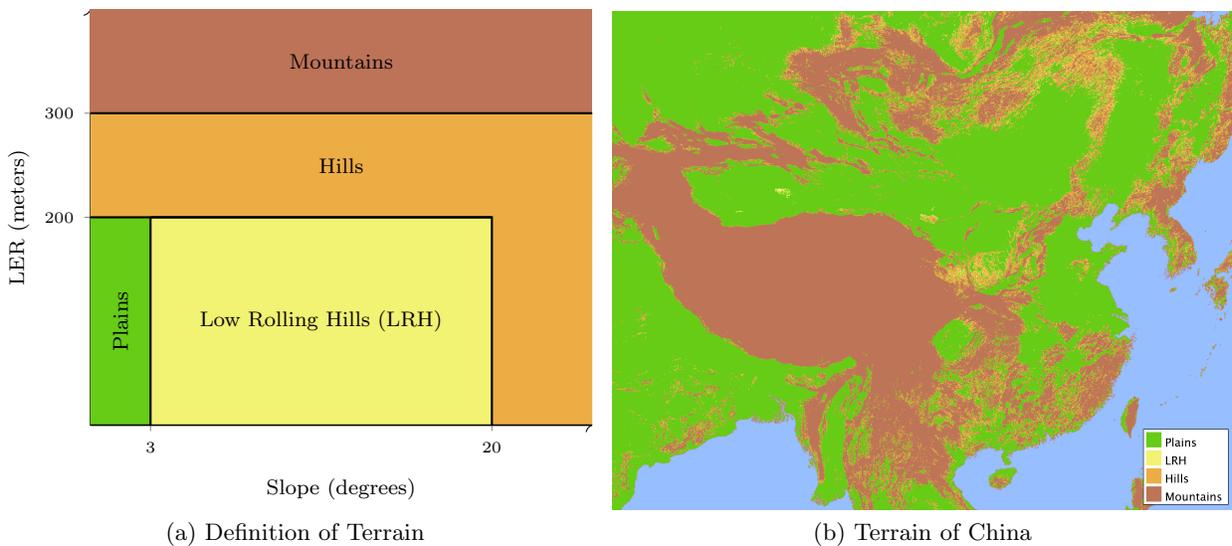


Figure B.3: Definition of Terrains, and the Terrain of China

Notes: Panel (a) defines the four types of terrains by slope and local elevation range (LER), conditional on elevation smaller than or equal to 2500 meters. All the areas with an elevation greater than 2500 meters are defined as “mountains”. Panel (b) shows the terrain of China following this definition.

Chapter 3.0.3 in the *Land Regulations* defines the terrains. Both plains and LRH are areas with a local elevation range (LER) of less than 200 meters; Plains are further defined as the areas where the slopes are smaller than 3 degrees, and low rolling hills as the areas where the slopes are between 3 and 20 degrees.

The definition of “hills” and “mountains” are less clear. the *Land Regulations* group the “hills” and the “mountains” into one category (Zhong Qiu Shan Ling) and defines them as the terrains with greater than 20 degrees of slope, or with LER greater than 200 meters. This aggregation of hills and mountains is because, by the time of the publication of the *Land Regulations* in 1999,

the Ministry of Transportation no longer distinguishes between these two types of terrains. The Ministry also publishes the official definition of the technical terms in the *Standard of Technical Terms for Highway Engineering*. Unfortunately, the distinction between “hills” and “mountains” are subjective and non-quantitative in this official publication. The “hills” are defined as the areas with relatively large LER but without characteristics of a mountain, such as a ridge, peak, or base. The “mountains”, on the other hand, are the terrains with large LER and observable characteristics. To quantify the design speed by terrain, we need a systematic way to distinguish between the hills and the mountains. The need to distinguish these two terrains arises because, in the 1988 revision of the standards, the design speed of highways differs between the hills and the mountains, as seen in Table 1.

To separately define hills and mountains, we turn to the official definition used by the UN Environment Programme World Conservation Monitoring Centre (UNEP-WCMC), which defines “mountains” as those with an LER greater than 300 meters or with an elevation greater than 2500 meters, following [Kapos et al. \(2000\)](#). Taking the *Land Regulations* and the UNEP definitions together, we define the “hills” as the terrains with more than 20 degrees of slope or with an LER between 200 and 300 meters, and the “mountains” as those with more than 300 meters in LER, or more than 2500 meters in elevation.

We use the data from GTOPO30 from the USGS for the slope and the LER data. The LER is then inferred from the elevation data. The slope and LER were evaluated for a five-pixel radius following [Kapos et al. \(2000\)](#).

B.3 Data Sources:

In addition to the maps and publications mentioned above, we also used the following datasets:

1. The **GTOPO30** from USGS. We use this data source to define the terrains as specified above. In particular, we use four sets of data to cover China: E060N40, E060N90, E100N40, and E100N90. This dataset provides the slope and elevation data with a grid spacing of 30 arc seconds, which is approximately 1 kilometer.
2. The **One Percent Population Survey** in 2005 and 2015. We use these two datasets to estimate the stay rates between 2000 and 2005, and between 2010 and 2015. In both years, we define the stay rate of a location as the fraction of individuals whose current location is the same as the self-reported location five years ago.

3. The **Custom Transaction Dataset** between 2000 and 2005. This dataset is used to identify the coastal prefectures that directly trade with the international markets.
4. The **City Statistical Yearbooks**. This dataset provides data on the goods and passenger transportation through each prefecture by mode of transportation. We drop air traffic in passenger transportation as it is only responsible for a small fraction of the total traffic.
5. The **Statistical Yearbooks** of various provinces. We used these data to construct the initial population and output in 1994. Many prefectures in our 291-prefecture sample were upgraded from “counties” after 1994 and were not covered by the City Statistical Yearbook back then. We reconstruct the population and output of these prefectures using the county-level data from the provincial-level Statistical Yearbooks.
6. The **Investment Climate Survey** conducted by the World Bank in 2002. We use the firm-level information to estimate the fraction of out-of-prefecture sales of Chinese firms, the same as in [Ma and Tang \(2020\)](#).
7. The **AIS Vessel Data**. This dataset reports the location, speed, and heading of vessels. We extract the speed of moving vessels, excluding fixed platforms, military vessels, and policy vessels. We use these data to approximate the average sailing speed.
8. The **Hukou Index** are constructed using Peking University Law Information Database (<http://www.lawinfochina.com/>). Following [Fan \(2019\)](#), we collect the laws and regulations implemented at the prefecture level that are potentially related to hukou reform by searching a set of keywords, including any combination of “hukou” or “huji” (also means hukou) with “gaige” (reform) or “guanli” (management), together with “chengshihua” or “chengzhenhua” (both mean urbanization) and “luohu” or “ruhu” (both mean granting hukou). The same scoring system (on a scale of 0-6 by adding up the subscores obtained from central districts and other parts of a city) and criteria (depending on housing tenure status and the length of contributing to local social security) are also applied as in [Fan \(2019\)](#) to our 291 prefectures.
9. The **Long Distance Passenger Bus Time Table** data comes from the digitized tables published by Tong Xin Publishers in *China Passenger Road Service Table Time*. The book covers the bus services originated from all the main prefecture cities in China in 1999. A typical entry covers an operating bus route and includes the following information: the origin and the final stations, distance in km, major stops en route, travel time between the origin

and the final station, ticket price, and vehicle types. As the data do not provide a detailed timetable between all the en route stops, our analysis focuses on the travel time between the origin and the final. We control the number of stops in the regression, as more stops usually lead to longer travel time.

The dataset used in our analysis is a subset of the book in which both the origin and the final stations are in our 291 prefecture cities. This sample selection criterion excludes the stations that originate or terminate in smaller counties outside of the major urban area of a prefecture. Our estimation sample contains 1,823 bus routes. These routes are representative, covering 257 out of 291 prefectures in our quantitative analysis. In cases where multiple services are available between a prefecture pair, we use the fastest service as a proxy for the travel time on the road network.

10. The **Railway Time Table** data come from www.12306.cn. We use data in year 2002 and 2010. Each entry in the dataset contains the detailed timetable of a train route, including the following variables: the name of the route, the list of en route stations (including the origin and the final stations), and the arrival and departure time at each station. Based on this information, one can compute the travel time between all station pairs on a route. Typically, if there are n stations on a route, we can extract up to $n(n-1)/2$ travel time records between stations. We then include all the station pairs in which both stations are located in the main urban area of a prefecture. In cases where multiple rail routes exist between a prefecture pair, we use the fastest service as a proxy for the travel time on the rail network. Our estimation sample is representative, as it covers 182 out of 291 prefectures in the main sample. Our railroad timetable data contain 3,412 observed travel times in a given year, out of which 534 pairs see changes in travel time between 2002 and 2010. As we only observe a subset of passenger travel times offered by the Ministry of Transportation, the travel time of a specific train route might not change despite the improvements in the overall railway network.

C Details of the Model

C.1 Dynamic Discrete Choice

Recall that the dynamic discrete choice problem defined in the main text, is replicated here for reference:

$$v_{jt} = \log \left(\phi_{jt} \frac{w_{jt}}{P_{jt}} \right) + \max_{i, r^m \in \cup_{m=1}^M \mathcal{R}_{ij}^m} \left\{ \delta V_{i,t+1} - \bar{\lambda}_{it} - \left(\sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} \right) + \kappa \cdot \varepsilon_{it, r^m} \right\}, \quad (\text{C.1})$$

Different from [Caliendo et al. \(2019\)](#), the individuals choose a destination i and a route to go from j to i among M types of transportation modes, $r^m \in \cup_{m=1}^M \mathcal{R}_{ij}^m$.

Denote $\zeta_{it, r^m} = \delta V_{i,t+1} - \bar{\lambda}_{it} - \left(\sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} \right) + \kappa \cdot \varepsilon_{it, r^m}$. Recall that ε_{it, r^m} follows a Generalized Type-I Extreme Value Distribution (GEV-I) with location parameter $\bar{\gamma}$ and a scale parameter of 1. As ζ_{it, r^m} is a linear transformation of ε_{it, r^m} , it is also a GEV-I with a location parameter $\left[\delta V_{i,t+1} - \bar{\lambda}_{it} - \left(\sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} \right) \right] + \kappa \bar{\gamma}$ and a scale parameter κ , because GEV-I is closed under linear transformations.

In light of this, the maximization problem in Equation (C.2) is equivalent to maximizing over countably many ζ_{it, r^m} with a common scale parameter κ and different location parameters:

$$v_{jt} = \log \left(\phi_{jt} \frac{w_{jt}}{P_{jt}} \right) + \max_{i, r^m \in \cup_{m=1}^M \mathcal{R}_{ij}^m} \{ \zeta_{it, r^m} \}. \quad (\text{C.2})$$

As GEV-I is also closed under maximization, it is straightforward to see that the distribution function of “ $\max_{i, r^m \in \cup_{m=1}^M \mathcal{R}_{ij}^m} \{ \zeta_{it, r^m} \}$ ” must be:

$$\begin{aligned} \Pr \left(\max_{i, r^m \in \cup_{m=1}^M \mathcal{R}_{ij}^m} \{ \zeta_{it, r^m} \} < x \right) &= \prod_{i=1}^J \prod_{r^m \in \cup_{m=1}^M \mathcal{R}_{ij}^m} \Pr (\zeta_{it, r^m} < x) \\ &= \prod_{i=1}^J \prod_{m=1}^M \prod_{r^m \in \mathcal{R}_{ij}^m} \exp \left(- \exp \left(- \frac{x - \left[\delta V_{i,t+1} - \bar{\lambda}_{it} - \left(\sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} \right) \right] - \kappa \bar{\gamma}}{\kappa} \right) \right) \\ &= \exp \left(- \sum_{i=1}^J \sum_{m=1}^M \sum_{r^m \in \mathcal{R}_{ij}^m} \exp \left(- \frac{x - \left[\delta V_{i,t+1} - \bar{\lambda}_{it} - \left(\sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} \right) \right] - \kappa \bar{\gamma}}{\kappa} \right) \right). \end{aligned}$$

Notice that:

$$\begin{aligned}
& \sum_{i=1}^J \sum_{m=1}^M \sum_{r^m \in \mathcal{R}_{ij}^m} \exp \left(- \frac{x - \left[\delta V_{i,t+1} - \bar{\lambda}_{it} - \left(\sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} \right) \right] - \kappa \bar{\gamma}}{\kappa} \right) \\
&= \exp \left\{ \log \left[\sum_{i=1}^J \sum_{m=1}^M \sum_{r^m \in \mathcal{R}_{ij}^m} \exp \left(- \left(\frac{x - \kappa \bar{\gamma}}{\kappa} - \frac{\left[\delta V_{i,t+1} - \bar{\lambda}_{it} - \left(\sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} \right) \right]}{\kappa} \right) \right) \right] \right\} \\
&= \exp \left\{ \log \left[\exp \left(- \frac{x - \kappa \bar{\gamma}}{\kappa} \right) \sum_{i=1}^J \sum_{m=1}^M \sum_{r^m \in \mathcal{R}_{ij}^m} \exp \left(\frac{\delta V_{i,t+1} - \bar{\lambda}_{it} - \left(\sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} \right)}{\kappa} \right) \right] \right\} \\
&= \exp \left\{ - \frac{x - \kappa \bar{\gamma}}{\kappa} + \log \left[\sum_{i=1}^J \sum_{m=1}^M \sum_{r^m \in \mathcal{R}_{ij}^m} \exp \left(\frac{\delta V_{i,t+1} - \bar{\lambda}_{it} - \left(\sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} \right)}{\kappa} \right) \right] \right\} \\
&= \exp \left\{ - \frac{x - \kappa \bar{\gamma} - \kappa \log \left[\sum_{i=1}^J \sum_{m=1}^M \sum_{r^m \in \mathcal{R}_{ij}^m} \exp \left(\frac{\delta V_{i,t+1} - \bar{\lambda}_{it} - \left(\sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} \right)}{\kappa} \right) \right]}{\kappa} \right\}.
\end{aligned}$$

Therefore, the distribution of $\max_{i,r^m \in \cup_{m=1}^M \mathcal{R}_{ij}^m} \{ \zeta_{it,r^m} \}$ is another GEV-I with the location parameter:

$$\kappa \log \left[\sum_{i=1}^J \sum_{m=1}^M \sum_{r^m \in \mathcal{R}_{ij}^m} \exp \left(\frac{\delta V_{i,t+1} - \bar{\lambda}_{it} - \left(\sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} \right)}{\kappa} \right) \right] + \kappa \bar{\gamma},$$

and the scale parameter κ . It directly follows from the property of GEV-I that:

$$\begin{aligned}
\mathbb{E} \left[\max_{i,r^m \in \cup_{m=1}^M \mathcal{R}_{ij}^m} \{ \zeta_{it,r^m} \} \right] &= \kappa \log \left[\sum_{i=1}^J \sum_{m=1}^M \sum_{r^m \in \mathcal{R}_{ij}^m} \exp \left(\frac{\delta V_{i,t+1} - \bar{\lambda}_{it} - \left(\sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} \right)}{\kappa} \right) \right] \\
&= \kappa \log \left[\sum_{i=1}^J \exp \left(\frac{\delta V_{i,t+1}}{\kappa} \right) \sum_{m=1}^M \sum_{r^m \in \mathcal{R}_{ij}^m} \exp \left(- \frac{\bar{\lambda}_{it} + \left(\sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} \right)}{\kappa} \right) \right].
\end{aligned}$$

Now, define λ_{ijt} as:

$$\begin{aligned}
\lambda_{ijt} &= -\kappa \log \left[\sum_{m=1}^M \sum_{r^m \in \mathcal{R}_{ij}^m} \exp \left(-\frac{\bar{\lambda}_{it} + \left(\sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} \right)}{\kappa} \right) \right] \\
&= -\kappa \log \left[\exp \left(-\frac{\bar{\lambda}_{it}}{\kappa} \right) \sum_{m=1}^M \sum_{r^m \in \mathcal{R}_{ij}^m} \exp \left(-\frac{\left(\sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} \right)}{\kappa} \right) \right] \\
&= \bar{\lambda}_{it} - \kappa \log \left[\sum_{m=1}^M \sum_{r^m \in \mathcal{R}_{ij}^m} \exp \left(-\frac{\left(\sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} \right)}{\kappa} \right) \right],
\end{aligned}$$

which is the expected migration costs across all the possible routes and modes that start from j , conditional on moving to location i , the same as in the main text in equation (4). We can then re-write the expectation term as:

$$\begin{aligned}
\mathbb{E} \left[\max_{i, r^m \in \cup_{m=1}^M \mathcal{R}_{ij}^m} \{ \zeta_{it, r^m} \} \right] &= \kappa \log \left[\sum_{i=1}^J \exp \left(\frac{\delta V_{i, t+1}}{\kappa} \right) \exp \left(-\frac{\lambda_{ijt}}{\kappa} \right) \right] \\
&= \kappa \log \left[\sum_{i=1}^J \exp \left(\frac{\delta V_{i, t+1} - \lambda_{ijt}}{\kappa} \right) \right],
\end{aligned}$$

and therefore, we arrive at equation (3) in the main text:

$$\mathbb{E}(v_{jt}) = V_{jt} = \log \left(\phi_{jt} \frac{w_{jt}}{P_{jt}} \right) + \kappa \log \left[\sum_{i=1}^J \exp \left(\frac{\delta V_{i, t+1} - \lambda_{ijt}}{\kappa} \right) \right].$$

C.2 Solution of λ_{ijt}

Starting from the definition of λ_{ijt} as in equation (4), we follow the steps in the Appendix D.1 of [Allen and Arkolakis \(2022\)](#) and extend their methods to allow for multiple modes of transportation. In particular, we enumerate the set \mathcal{R}_{ij}^m by the number of steps, K . Denote the set of all routes

from j to i under mode m with length K as \mathcal{R}_{ij}^{mK} , we first re-write λ_{ijt} as:

$$\begin{aligned}\lambda_{ijt} &= \bar{\lambda}_{it} - \kappa \log \left[\sum_{m=1}^M \sum_{r^m \in \mathcal{R}_{ij}^m} \exp \left(- \frac{\left(\sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} \right)}{\kappa} \right) \right] \\ &= \bar{\lambda}_{it} - \kappa \log \left[\sum_{m=1}^M \sum_{K=0}^{\infty} \sum_{r^m \in \mathcal{R}_{ij}^{mK}} \exp \left(- \frac{1}{\kappa} \sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} \right) \right] \\ &= \bar{\lambda}_{it} - \kappa \log \left\{ \sum_{m=1}^M \sum_{K=0}^{\infty} \sum_{r^m \in \mathcal{R}_{ij}^{mK}} \left[\prod_{k=1}^K \exp \left(- \frac{d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}}}{\kappa} \right) \right] \right\}.\end{aligned}$$

Note that the last step converted the additive transportation costs in $d_{(\cdot)}^{m\mathbb{P}}/\kappa$ into multiplicative costs in the unit of $\exp\left(\frac{d_{(\cdot)}^{m\mathbb{P}}}{\kappa}\right)$, similar to the main text in [Allen and Arkolakis \(2022\)](#). To proceed, we denote the adjacency matrix as $\mathbf{F}_t^{m\mathbb{P}}$, where $F_{ijt}^{m\mathbb{P}} = \exp\left(-\frac{d_{ijt}^{m\mathbb{P}}}{\kappa}\right)$ is the (i, j) th element of the matrix, and $(\mathbf{F}_t^{m\mathbb{P}})^K$ is the matrix raised to the power K . Note that we can enumerate all the routes with length K :

$$\begin{aligned}\lambda_{ijt} &= \bar{\lambda}_{it} - \kappa \log \left\{ \sum_{m=1}^M \sum_{K=0}^{\infty} \left[\sum_{k_1=1}^J \sum_{k_2=1}^J \cdots \sum_{k_{K-1}=1}^J \left(F_{i,k_1,t}^{m\mathbb{P}} \times F_{i,k_2,t}^{m\mathbb{P}} \cdots \times F_{k_{K-1},j,t}^{m\mathbb{P}} \right) \right] \right\} \\ &= \bar{\lambda}_{it} - \kappa \log \left\{ \sum_{m=1}^M \sum_{K=0}^{\infty} (\mathbf{F}_t^{m\mathbb{P}})^K_{ij} \right\},\end{aligned}$$

where $(\mathbf{F}_t^{m\mathbb{P}})^K_{ij}$ is the (i, j) th element of the adjacency matrix raised to the power K . Define $\mathbf{B}_t^{m\mathbb{P}}$ as the Leontief Inverse of $\mathbf{F}_t^{m\mathbb{P}}$:

$$\mathbf{B}_t^{m\mathbb{P}} \equiv (\mathbf{I} - \mathbf{F}_t^{m\mathbb{P}})^{-1} = \sum_{K=0}^{\infty} (\mathbf{F}_t^{m\mathbb{P}})^K.$$

We then arrive at the solution of λ_{ijt} as shown in the main text:

$$\lambda_{ijt} = \bar{\lambda}_{it} - \kappa \log \left(\sum_{m=1}^M b_{ijt}^{m\mathbb{P}} \right),$$

where $b_{ijt}^{m\mathbb{P}}$ is the (i, j) th element of the matrix $\mathbf{B}_t^{m\mathbb{P}}$.

C.3 Solution of τ_{ijt}

Similar to the solution of λ_{ijt} , we start from equation (9) in the main text and enumerate all the possible routes under mode m by length, we can re-write the above equation as:

$$\tau_{ij}^{-\theta} = (\bar{\tau})^{-\theta} \sum_{m=1}^M \sum_{K=1}^{\infty} \left(\mathbf{F}_{ijt}^{mf} \right)^K,$$

where $\mathbf{F}_t^{mf} \equiv \left[\left(d_{ijt}^{mf} \right)^{-\theta} \right]$ is the adjacency matrix of mode m , F_{ijt}^{mf} is the (i, j) th element of the matrix, and $\left(\mathbf{F}_t^{mf} \right)^K$ is the matrix raised to the power K . The power sequence can be evaluated as follows:

$$\sum_{K=1}^{\infty} \left(\mathbf{F}_t^{mf} \right)^K = \left(\mathbf{I} - \mathbf{F}_t^{mf} \right)^{-1} \equiv \mathbf{B}_t^{mf}.$$

Denote the (i, j) th element of the matrix \mathbf{B}_t^{mf} as b_{ijt}^{mf} , we can then write τ_{ijt} as:

$$\tau_{ijt} = \bar{\tau} \cdot \left(\sum_{m=1}^M b_{ijt}^{mf} \right)^{-\frac{1}{\theta}}.$$

C.4 Solving the Model in Levels

Conditional on observing $\{L_{j0}\}$, Θ , and a finite sequence of $\Theta_t, t = 1, 2, \dots, T$, we solve the model in levels as follows.

1. Assume that the fundamentals are constants after period T , so that $\Theta_{T+k} = \Theta_t, \forall k > 0$.
2. Solve the stationary equilibrium after T , denoted as $\bar{\Upsilon} = \{w_j, L_j, P_j, V_j, \mu\}$ from the following

system of equations:

$$w_j L_j = \sum_{i=0}^J \frac{(w_j \tau_{ij})^{-\theta} (\bar{A}_j L_j^\alpha)^\theta}{\sum_{k=1}^J (w_k \tau_{ik})^{-\theta} (A_k L_k^\alpha)^\theta} w_i L_i \quad (\text{C.3})$$

$$P_j = \Gamma(1 + (1 - \eta/\theta)) \left(\sum_{i=1}^J (w_i \tau_{ji})^{-\theta} (A_i L_i^\alpha)^\theta \right)^{-1/\theta} \quad (\text{C.4})$$

$$V_j = \log \left(\bar{\phi}_j L_j^\beta \frac{w_j}{P_j} \right) + \kappa \log \left(\sum_{i=1}^J \exp(\delta V_i - \lambda_{ij})^{1/\kappa} \right) \quad (\text{C.5})$$

$$\mu_{ij} = \frac{\exp(\delta V_i - \lambda_{ij})^{1/\kappa}}{\sum_{i'=1}^J \exp(\delta V_{i'} - \lambda_{i'j})^{1/\kappa}} \quad (\text{C.6})$$

$$L_i = \sum_{j=1}^J \mu_{ij} L_j \quad (\text{C.7})$$

To solve this system:

- (a) Guess $\{L_j\}$.
- (b) Solve $\{w_j\}$ from equation (C.3).
- (c) Compute $\{P_j\}$ using $\{w_j, L_j\}$ and equation (C.4).
- (d) Solve $\{V_j\}$ from equation (C.5). Note that in order to efficiently compute V_j , re-write equation (C.5) as:

$$\begin{aligned} V_j &= \log \left(\bar{\phi}_j L_j^\beta \frac{w_j}{P_j} \right) + \kappa \log \left(\sum_{i=1}^J \exp \left(\frac{\delta V_i}{\kappa} - \frac{\lambda_{ij}}{\kappa} \right) \right) \\ &= \log \left(\bar{\phi}_j L_j^\beta \frac{w_j}{P_j} \right) + \kappa \log \left(\sum_{i=1}^J \exp \left(\frac{\delta V_i}{\kappa} \right) \times \exp \left(-\frac{\lambda_{ij}}{\kappa} \right) \right) \end{aligned}$$

In this way, we can evaluate the sum inside the second log function as the inner product between two vectors, $\left\{ \exp \left(\frac{\delta V_i}{\kappa} \right) \right\}_{i=1}^J$ and $\left\{ \exp \left(-\frac{\lambda_{ij}}{\kappa} \right) \right\}$. Importantly, the second vector is not a function of V_i and, therefore, can be evaluated outside the loop.

- (e) Compute μ from equation (C.6).
 - (f) Update $\{L_j\}$ from equation (C.7). Repeat until converge.
3. Pick a $T' > T$ and guess a sequence of $\{V_{jt}\}_{t=1}^{T'}$, where $V_{jT'} = V_j$ as computed from above.
 4. With the $\{V_{jt}\}_{t=1}^{T'}$ and the initial population, we can compute a sequence of $\{\mu_{jt}\}_{t=1}^{T'}$ and $\{L_{jt}\}_{t=1}^{T'}$ using equations (5) and (6) forward from $t = 1, \dots, T'$. In particular, re-write

equation (5) as:

$$\mu_{ijt} = \frac{\exp\left(\frac{\delta V_{i,t+1}}{\kappa}\right) \exp\left(-\frac{\lambda_{ijt}}{\kappa}\right)}{\sum_{i'=1}^J \exp\left(\frac{\delta V_{i',t+1}}{\kappa}\right) \exp\left(-\frac{\lambda_{i'jt}}{\kappa}\right)}$$

5. With the sequence of $\{L_{jt}\}_{t=1}^{T'}$, solve the sequence of $\{w_{jt}\}_{t=1}^{T'}$ using equation (11).

6. Compute the sequence of $\{P_{jt}\}_{t=1}^{T'}$ using equation (C.8):

$$P_{jt} = \Gamma\left(\frac{\theta + 1 - \eta}{\theta}\right) \left(\sum_{i=0}^J (w_{it}\tau_{jit})^{-\theta} (A_{it})^\theta\right)^{-1/\theta}, \quad (\text{C.8})$$

where $\Gamma(\cdot)$ is the standard Gamma function.

7. Update the sequence of $\{V_{jt}\}_{t=1}^{T'}$ backwards from $T' - 1$ to $t = 1$ using equation (3):

$$V_{jt} = \log\left(\phi_{jt} \frac{w_{jt}}{P_{jt}}\right) + \kappa \log\left(\sum_{i=1}^J \exp\left(\frac{\delta V_{i,t+1}}{\kappa}\right) \exp\left(-\frac{\lambda_{ijt}}{\kappa}\right)\right).$$

8. Repeat until convergence.

D Quantification

D.1 Moment Conditions

Trade Shares by Mode The volume of sales from j to i via mode m is:

$$\sum_{r^m \in \mathcal{R}_{ij}^m} \pi_{ijt,r^m} X_{ij} = \frac{(w_{jt}/A_{jt})^{-\theta} \sum_{K=1}^{\infty} (\mathbf{F}_t^{mf})^K}{\sum_{k=0}^J (w_{kt}/A_{kt})^{-\theta} \tau_{ikt}^{-\theta}} X_{ij} = \frac{(w_{jt}/A_{jt})^{-\theta} b_{ijt}^{mf}}{\sum_{k=0}^J (w_{kt}/A_{kt})^{-\theta} \tau_{ikt}^{-\theta}} X_{ij},$$

The share of sales from j to i via mode m can be written as:

$$s_{ijt}^{mf} = \frac{(w_{jt}/A_{jt})^{-\theta} b_{ijt}^m}{\sum_{k=0}^J (w_{kt}/A_{kt})^{-\theta} \tau_{ikt}^{-\theta}} X_{ij} = \frac{b_{ijt}^{mf}}{\sum_{m'=1}^M \frac{(w_{jt}/A_{jt})^{-\theta} b_{ijt}^{m'}}{\sum_{k=0}^J (w_{kt}/A_{kt})^{-\theta} \tau_{ikt}^{-\theta}} X_{ij}},$$

and the share of all the sales from j via mode m , denoted as s_{jt}^{mf} is the weighted average of s_{ijt}^{mf} across destinations excluding ROW, where the weight is X_{ijt}/X_{jt} :

$$\begin{aligned} s_{jt}^{mf} &= \sum_{i=1}^J s_{ijt}^{mf} \frac{X_{ijt}}{X_{jt}} \\ &= \sum_{i=1}^J \left(\frac{b_{ijt}^{mf}}{\sum_{m'=1}^M b_{ijt}^{m'f}} \right) \frac{X_{ijt}}{X_{jt}}. \end{aligned}$$

Migration Shares by Mode The fraction of migrants from j to i via mode m at time t is:

$$\begin{aligned} s_{ijt}^{m\mathbb{P}} &= \frac{\sum_{r^m \in \mathcal{R}_{ij}^m} \exp \left(\frac{\delta V_{i,t+1} - \bar{\lambda}_{it} - \left(\sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} \right)}{\kappa} \right)}{\sum_{m'=1}^M \sum_{r^{m'} \in \mathcal{R}_{ij}^{m'}} \exp \left(\frac{\delta V_{i,t+1} - \bar{\lambda}_{it} - \left(\sum_{k=1}^K d_{r_{k-1}^{m'}, r_k^{m'}, t}^{m'\mathbb{P}} \right)}{\kappa} \right)} \\ &= \frac{\sum_{r^m \in \mathcal{R}_{ij}^m} \exp \left(\frac{- \left(\sum_{k=1}^K d_{r_{k-1}^m, r_k^m, t}^{m\mathbb{P}} \right)}{\kappa} \right)}{\sum_{m'=1}^M \sum_{r^{m'} \in \mathcal{R}_{ij}^{m'}} \exp \left(\frac{- \left(\sum_{k=1}^K d_{r_{k-1}^{m'}, r_k^{m'}, t}^{m'\mathbb{P}} \right)}{\kappa} \right)} = \frac{\sum_{K=0}^{\infty} (\mathbf{F}_t^{m\mathbb{P}})_{ij}^K}{\sum_{m'=1}^M (\mathbf{F}_t^{m'\mathbb{P}})_{ij}^K} = \frac{b_{ijt}^{m\mathbb{P}}}{\sum_{m'=1}^M b_{ijt}^{m'\mathbb{P}}}. \end{aligned}$$

Therefore the total out-migration from j to all the other (non-ROW) locations via mode m is the weighted average of the above probability, where the weight is the migration flow:

$$s_{jt}^{m\mathbb{P}} = \sum_{i=1}^J \left(\frac{b_{ijt}^{m\mathbb{P}}}{\sum_{m'=1}^M b_{ijt}^{m'\mathbb{P}}} \right) \mu_{ijt} = \sum_{i=1}^J \left(\frac{b_{ijt}^{m\mathbb{P}}}{\sum_{m'=1}^M b_{ijt}^{m'\mathbb{P}}} \right) \frac{L_{ijt}}{L_{jt}}.$$

Note that the passenger volume data in *City Statistical Yearbook* cover migrants and short-term travelers for business and leisure purposes. By matching $s_{jt}^{m\mathbb{P}}$ to the observed shares, we implicitly assume that there are no systematic differences in the choice of travel modes between migrants and short-term travelers.

Coefficient of Variations The model counterparts of the CV of total freight and passenger volume in location j , time t are, respectively:

$$CV_{jt}^f = \frac{\text{std}(\{X_{ijt}\}_{i>0 \cap i \neq j})}{\sum_{i>0 \cap i \neq j} X_{ijt} / (J-1)}, \quad CV_{jt}^p = \frac{\text{std}(\{L_{ijt}\}_{i>0 \cap i \neq j})}{\sum_{i>0 \cap i \neq j} L_{ijt} / (J-1)},$$

where $\text{std}(\cdot)$ computes the standard deviation of a vector. The moment conditions are the vector of the average CV across the prefecture in each year between 1994 and 2017: $\left\{ \frac{\sum_{j=1}^J \text{CV}_{jt}^{\text{f}}}{J} \right\}_{t=1}^{23}$ and $\left\{ \frac{\sum_{j=1}^J \text{CV}_{jt}^{\text{p}}}{J} \right\}_{t=1}^{23}$.

D.2 Estimation

Estimation In summary, denote the moment conditions in the data as the vector $\bar{\mathcal{S}}$, and the counter-parts in the model as

$$\mathcal{S}(\Theta^x) = \left\{ \frac{\sum_{j=1}^J s_{jt}^{m\text{f}}}{J}, \frac{\sum_{j=1}^J s_{jt}^{m\text{p}}}{J}, \frac{\sum_{j=1}^J \text{CV}_{jt}^{\text{f}}}{J}, \frac{\sum_{j=1}^J \text{CV}_{jt}^{\text{p}}}{J} \right\}_{t=1}^{23}.$$

Our estimation strategy is to find Θ^x to minimize the Euclidean distance between the data and the model moments:

$$\min_{\Theta^x} [\bar{\mathcal{S}} - \mathcal{S}(\Theta^x)] \mathbf{W} [\bar{\mathcal{S}} - \mathcal{S}(\Theta^x)]', \quad (\text{D.1})$$

where \mathbf{W} is the optimal weighting matrix. In this context, the optimal weighting matrix is the inverse of the variance-covariance matrix of the data moments $\bar{\mathcal{S}}$, computed by bootstrapping the data at the prefecture-mode-year level 1,000 times.

We use a combination of numerical optimization procedures to minimize the objective function. We start the optimization process using Particle Swarm Optimization with 48 particles. Upon convergence, we switch to grid search and pattern search to refine the results. In the end, we turn to gradient-based methods to speed up the convergence. The optimization procedure described above does not necessarily lead to a global minimum. However, the objective function is well-behaved in a wide neighborhood around our point estimate, as shown in Figure D.1.

Asymptotic Standard Errors We estimate the asymptotic standard errors following [Cameron and Trivedi \(2005\)](#). In particular, we use:

$$\widehat{\text{var}}(\Theta^x) = \left(\widehat{\mathbf{G}}' \mathbf{W} \widehat{\mathbf{G}} \right)^{-1} \widehat{\mathbf{G}}' \mathbf{W} \widehat{\Sigma} \mathbf{W} \widehat{\mathbf{G}} \left(\widehat{\mathbf{G}}' \mathbf{W} \widehat{\mathbf{G}} \right)^{-1}, \quad (\text{D.2})$$

In the equation above, $\widehat{\mathbf{G}}$ is an estimate of the gradient matrix, in which the i th row and the j th column is the partial derivative of the i th moment with respect to the j th parameter, evaluated

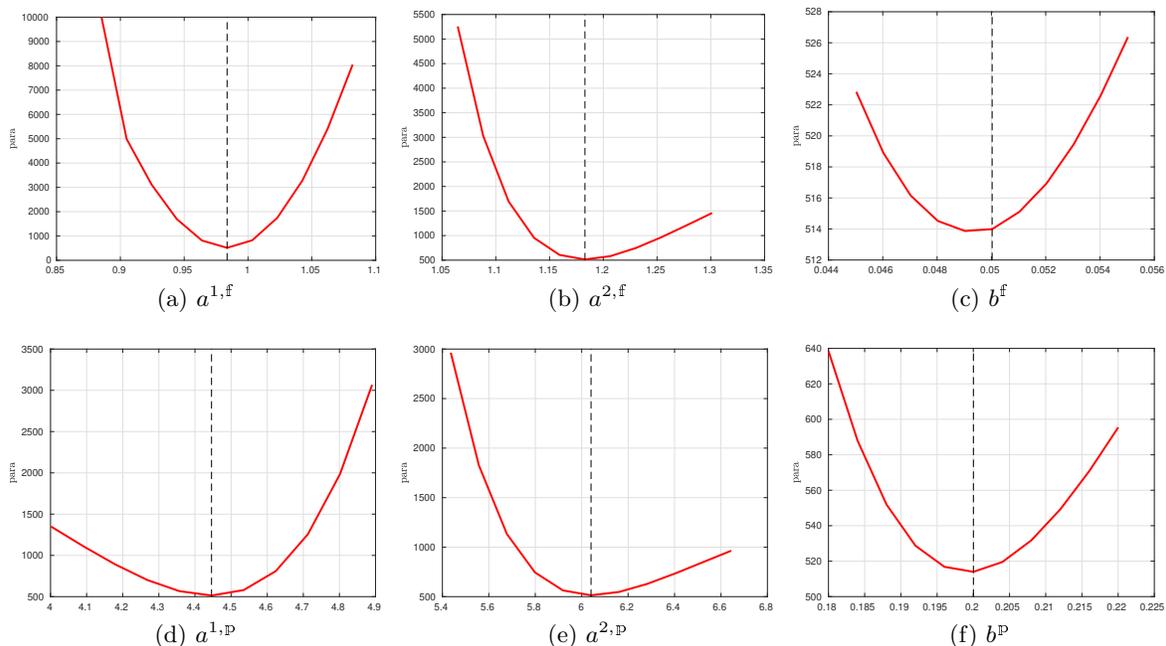


Figure D.1: The Local Grid Search around the Point Estimate

Notes: The panels in this figure plot the objective function in equation (D.1) against each parameter while holding the other parameters at the estimated value shown in Table 5. The dashed line indicates the estimated values of the parameter in question.

at Θ^x . We numerically approximate the gradient matrix. \mathbf{W} is the weighting matrix, and $\hat{\Sigma}$ is an estimate of the variance-covariance matrix of the moment conditions.

We estimate $\hat{\Sigma}$ using 1,000 bootstraps of the data moments, and set $\mathbf{W} = (\hat{\Sigma})^{-1}$. With the optimal weighting matrix, equation (D.2) simplifies to:

$$\widehat{\text{var}}(\Theta^x) = \left(\hat{\mathbf{G}}' \hat{\Sigma}^{-1} \hat{\mathbf{G}} \right)^{-1}, \quad (\text{D.3})$$

which we use to estimate the standard errors as reported in Table 5.

Identification This section provides more details related to the parameter identification in the GMM. As is common in structural estimation, the parameters are jointly identified by all the moment conditions. Nevertheless, we can still map the identification of each parameter to specific groups of moment conditions as suggested by the gradient matrix illustrated in Figure D.2. Recall that the gradient matrix, $\hat{\mathbf{G}}$, contains the partial derivative of the i th moment with respect to the j th parameter in its i th row and j th column. For ease of exposition, we present this 81-by-6 gradient matrix using a heatmap for each parameter in each panel of Figure D.2. In that figure,

the numbers are the partial derivatives of each moment to a parameter in question. For example, Panel (a) of the figure plots the column of 81 partial derivatives with respect to $a^{1,\text{f}}$, the costs related to freight transportation on the road networks. We break down the 81 partial derivatives into three groups: 1) the average road usage shares in each year, $\sum_{j=1}^J s_{jt}^{1,\text{f}}/J$ and $\sum_{j=1}^J s_{jt}^{1,\text{P}}/J$, 2) the average rail usage shares in each year, $\sum_{j=1}^J s_{jt}^{2,\text{f}}/J$ and $\sum_{j=1}^J s_{jt}^{2,\text{P}}/J$, and 3) the coefficients of variation of total freight and passenger traffic each year, $\text{CV}_{jt}^{\text{f}}$ and $\text{CV}_{jt}^{\text{P}}$. Note that we do not have data on passenger traffic usage in 2014.

The parameters of mode-specific usage costs are identified mainly by the usage shares in the data. For example, as suggested by Panel (a) of Figure D.2, a higher $a^{1,\text{f}}$ decreases the road usage share and increases the rail usage share of freight transportation, while those of the passenger transportation hardly responses to the changes in $a^{1,\text{f}}$. A higher cost of railway usage, $a^{2,\text{f}}$, moves the moment conditions in opposite directions. Therefore, the variations in road and rail usage share in the data for freight transportation helps to pin down both $a^{1,\text{f}}$ and $a^{2,\text{f}}$. Similarly, as suggested by Panels (c) and (d) of the same figure, the variations in road and rail usage shares for passenger transportation determine the costs of mode-specific costs of passenger transportation, $a^{1,\text{P}}$ and $a^{2,\text{P}}$.

The elasticity parameters of trade and migration with respect to travel time, h^{f} and h^{P} , are mostly determined by the coefficients of variation, as shown in the last two panels of Figure D.2. In both cases, higher elasticity increases the coefficient of variations of both trade and migration flows. Intuitively, higher elasticity amplifies the variations in travel time, so a slight change in T leads to significant variations in τ or λ . As a result, the variations in trade and migration flows also increase. In the other direction, as the elasticities converge to zero, the variations in travel time become irrelevant, and the τ and λ matrices converge to constants. In the limit the trade costs among locations are the same, as do the migration costs net of entry barriers. As a result, the variations in the observed trade and migration flows also decline.

Lastly, the moments on coefficients of variation are affected by changes in almost all parameters. This differs from the road and rail usage moments that display a stark partition between freight and passenger transportation: the moments related to freight transportation do not usually respond to the parameter related to passenger transportation, and vice versa. The coefficients of variation are highly sensitive due to the interconnected nature of trade and migration flows in the general equilibrium. For instance, an increase in population inflow to a location raises its demand, resulting in higher trade flows. Similarly, improved market access in a location increases its attractiveness to migrants, leading to higher migration flow. Consequently, any parameter that impacts the

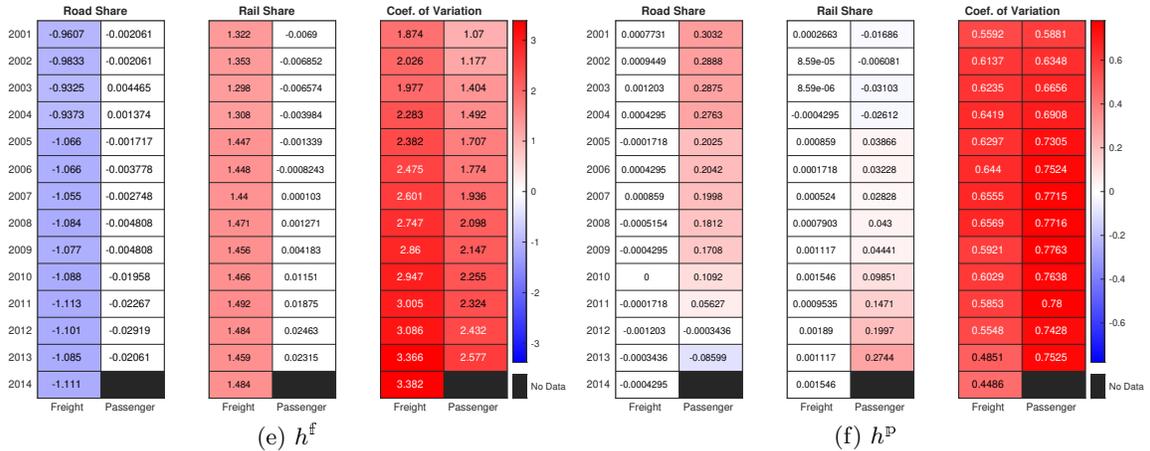
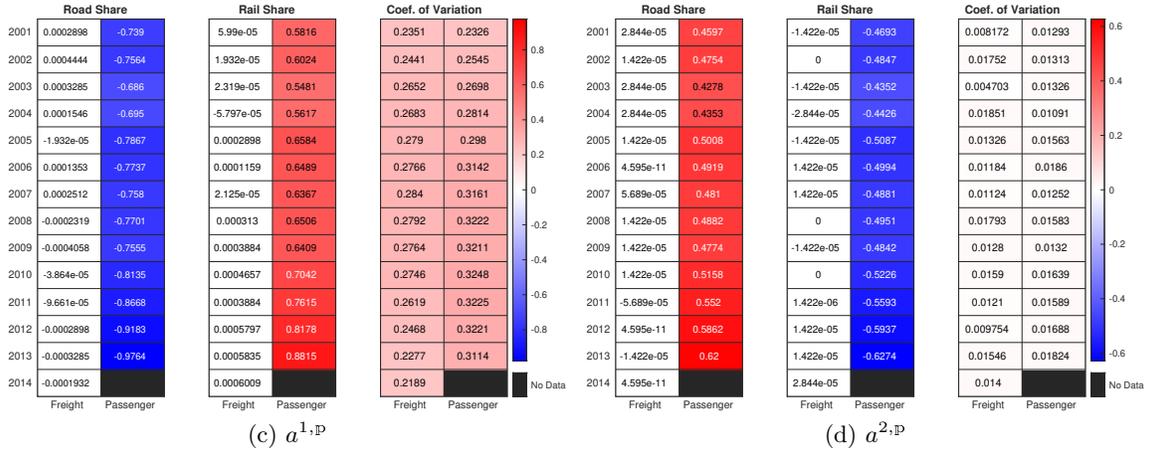
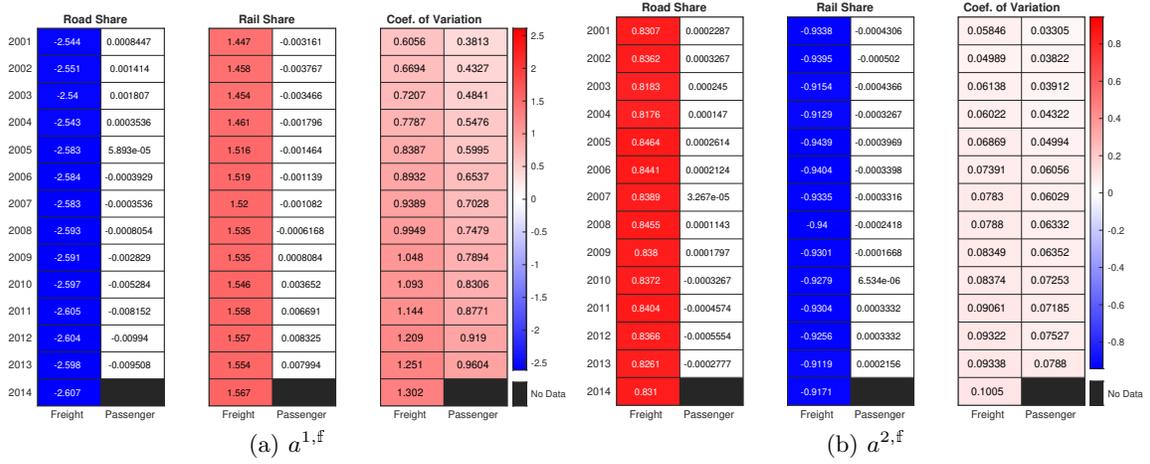


Figure D.2: The Gradient Matrix

Notes: This figure presents the estimated gradient matrix, \hat{G} , around the baseline estimates. Each panel reports the partial derivatives of the 81-moment conditions with respect to the parameter in the panel title.

distribution of one type of flow also influences the other.

D.3 Return to Investment in Transportation

We estimate the costs of transportation networks as follows. Wu et al. (2021) estimated that on average, China spent 8.9% of GDP on infrastructure between 1998 and 2007. Their definition of infrastructure includes three categories: 1) electricity, gas, and water, 2) transportation, and 3) information transmission, computer services, and software. *China Statistical Yearbooks* reported that between these years, on average, investments in the transportation sector account for 58% of the total investments in the three categories combined. Based on this information, we assume that, on average, China invests $8.9 * 0.58 = 5.16$ percent of GDP on transportation infrastructure each year. Using the real GDP data of China from Penn World Table (rgdpna in 2017 USD), we estimate that a simple sum over the investment in transportation infrastructure in each year between 1995 and 2017 is \$12.9 trillion. Converting the annual flow of investment to present values in 1995 using a discount rate of $\delta = 0.97$, the total costs of investment equals \$8.6 trillion, which is 184% of GDP in China in 1995.

To estimate the benefit of investments, we combine the counterfactual growth rates from our simulation with the observed GDP data. The accumulated growth rate of aggregate output in year t is denoted as $g_t = Y_t/Y_1 - 1$. The impact of infrastructure on growth rate in our model is then summarized as $\hat{g}_t = g_t/g'_t$, where g'_t is the accumulated growth rate when there is no expansion in the transportation network (the “no change” counterfactual), and g_t is the growth rate in the baseline equilibrium. For example, Figure 6 shows that by 2044, the changes in transportation infrastructure increase the accumulated growth rate from $g'_t = 4.7$ percent to $g_t = 7.4$ percent, leading to a gain of $\hat{g}_t = g_t/g'_t = 1.57$, or 57 percent. Denote the sequence of accumulated growth rate in the data as $\{g_t^{\text{data}}\}$. Our counter-factual accumulated growth rate is then defined as $\{g_t^{\text{counter}} = g_t^{\text{data}}/\hat{g}_t\}$. Starting from the observed GDP in 1995, we can compute a counterfactual sequence of GDP using the sequence of g_t^{counter} to proxy the world without investments in transportation infrastructure. To project the GDP growth between 2018 and 2044, we use the average annual growth rate in the data (5.76%) and the counterfactual (3.4%) during the sample period. We claim the difference between the observed and the counterfactual GDP sequence as the return to investments in transportation infrastructure. We similarly convert the flow of return to present values in 1995 using the same discount factor of δ .

The above exercise shows that the present value of the benefits between 1995 and 2017 is \$16.5

trillion. Compared to the present value of costs at \$8.6 trillion, the total return is $16.5/8.6 - 1 = 92.2\%$ or an annualized return rate of $1.92^{(1/22)} - 1 = 3.0\%$. Extending the accounting of benefits to 2044 yields a present value of the total benefits to be \$201 trillion. This implies that the accumulated return of transportation investment over 50 years is 2242% or an annualized return of 6.8%. This finding is similar to those in [Wu et al. \(2021\)](#) that estimated the return to overall infrastructure using firm-level data. They find that the return on infrastructure investment is 6 percent.

D.4 The Dynamic Response of Migration and Trade

As shown in [Figure 6](#), the short-run return to migration liberalization could be negative. For example, before 2010, the aggregate welfare under the λ -only simulation was often *worse* than the no-change case, suggesting a negative short-term return to migration liberalization. In the long run, the return to the reduction in λ starts to turn positive. The delayed return to migration liberalization comes from the forward-looking behavior of the migrants. If people expect a better passenger travel network and looser hukou restrictions in the future, they might delay the intended move to the more productive prefectures. [Figure D.3](#) confirms the dynamic impacts on population movements by plotting the evolution of the population in the most and least productive prefectures. As individuals expect lower migration costs in the future, the migration towards the most productive prefectures is delayed, as shown in [Panel \(a\)](#). Instead, people prolong their stay in less productive locations, as shown in [Panel \(b\)](#). Both forces lead to lower economic output in the short run. In the long run, however, the return to migration liberalization is substantial and long-lasting, facilitating population movements toward more productive locations. [Figure D.3](#) also highlights that the policy barriers ($\bar{\lambda}_{jt}$) in general exert stronger impacts on migration decisions than the passenger travel network ($\mathbf{T}_t^{m\mathbb{P}}$), a finding that echos those in [Li and Ma \(2022\)](#).

The return to trade liberalization, on the other hand, is always positive along the transition path. The immediate return to trade liberalization is straightforward because the trade decisions are static in the model, as seen in [equation \(11\)](#). In other words, the firms and individuals immediately respond to the changes in the trade network without any forward-looking concerns.

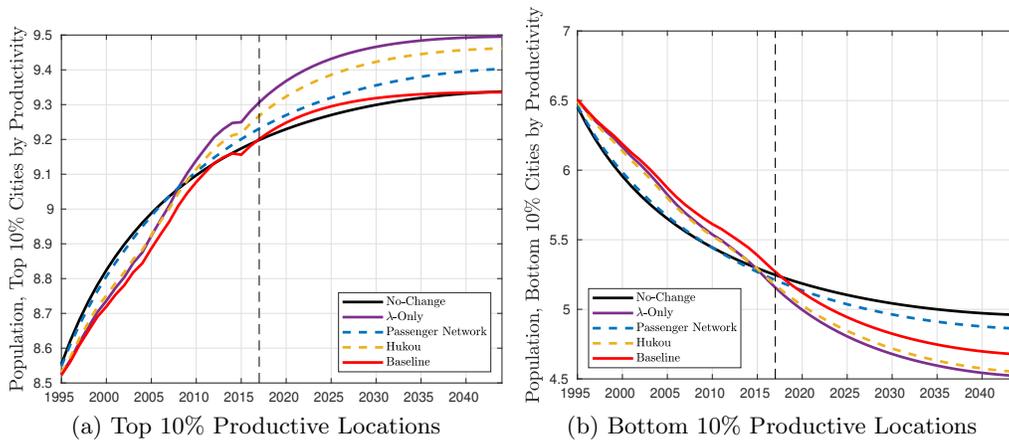


Figure D.3: The Dynamic Response to Migration Liberalization

Notes: This figure plots the transition paths of population flows across simulations. “No-Change” is the case in which both τ and λ are fixed to their 1995 levels. “λ-only” is the case in which τ is fixed at the 1995 level, and λ changes over time. “Passenger Network” is the case in which only \mathbf{T}_t^{mp} changes over time, “Hukou” is the case in which only $\bar{\lambda}_{jt}$ changes over time, and “Baseline” is the case in which both τ and λ changes over time as in the data. The y-axis is the population measured as the population share in percentage points, where the total population in China is 100. The vertical dashed line indicates 2017, the last year we have data on transportation networks.

E Model with Capital Accumulation

E.1 Model

The baseline model can be readily extended to include physical capital accumulation, which better mimics the phenomenal economic growth taking place in China during the sample period. Following [Kleinman et al. \(2023\)](#), in each location we introduce immobile landlords, who make consumption and investment decisions to maximize their lifetime utility:

$$v_{it}^k = \sum_{s=0}^{\infty} (\delta)^{t+s} \ln c_{it+s}^k,$$

where δ is the same discount factor as in the main model. The landlord's budget constraint is expressed as:

$$r_{it}k_{it} = P_{it}(c_{it}^k + k_{it+1} - (1 - \xi)k_{it}),$$

where r_{it} is the rate of return on capital and p_{it} is the price index in location i at time t . The parameter ξ is the capital depreciation rate. Following [Kleinman et al. \(2023\)](#), we assume that landlords can only invest locally, and capital is immobile across locations.

The instantaneous logarithm utility function enables a closed-form solution for the capital accumulation process:

$$k_{it+1} = \delta \left(1 - \xi + \frac{r_{it}}{P_{it}} \right) k_{it}. \tag{E.1}$$

Note that since both the price index and the rate of return to capital are endogenous and location-specific, and thus capital is accumulated at different speeds among locations despite the immobility of capital.

The production function for every variety ω now becomes

$$q_{jt}(\omega) = A_{jt} \left(\frac{\ell_{jt}}{\rho} \right)^{\rho} \left(\frac{k_{jt}}{1 - \rho} \right)^{1 - \rho}.$$

Zero profit condition implies

$$\frac{w_{jt}L_{jt}}{r_{jt}K_{jt}} = \frac{\rho}{1 - \rho}, \tag{E.2}$$

and in particular,

$$X_{jt} = w_{jt}L_{jt} + r_{jt}K_{jt} = \frac{1}{\rho}w_{jt}L_{jt}. \quad (\text{E.3})$$

The condition above implies that the trade balance condition described in Equation (11) still holds in the current scenario with capital accumulation. With a solution of w_{jt} , conditional on (L_{jt}, K_{jt}) , the interest rate is simply

$$r_{jt} = \frac{1 - \rho}{\rho} \frac{w_{jt}L_{jt}}{K_{jt}} \quad (\text{E.4})$$

The sequential competitive equilibrium with capital accumulation now consists of a sequence of endogenous variables $\Upsilon_t = \{w_{jt}, L_{jt}, V_{jt}, \mu_{ijt}, r_{jt}, k_{jt}\}$, conditional on time-invariant fundamentals $\bar{\Omega} = \{\bar{A}_j, \bar{\phi}_j\}$ and time-variant fundamentals $\Omega_t = \{\bar{\lambda}_{jt}, \mathbf{D}_t^{m\mathbb{P}}, \mathbf{D}_t^{mf}\}$, such that

1. Individuals maximize their life-time utility (3) by choosing a sequence of locations so that equations (5) and (6) hold.
2. Landlords maximize their lifetime utility by making the investment decision so that equation (E.1) holds.
3. Firms maximize their profits in each period and trade balance, so that equation (11) and (E.2) hold.

E.2 Solving the Model in Levels

Steady State Solving the stationary equilibrium requires a similar system of equations as those in (C.3) to (C.7) in Appendix C.4. To simplify notation, we denote the marginal costs of production as v_{jt} :

$$v_{jt} = (w_{jt})^\rho (r_{jt})^{1-\rho} \quad (\text{E.5})$$

With this notation, the system of equations that defines the steady state is as follows:

$$v_j = (w_j)^\rho (r_j)^{1-\rho} \quad (\text{E.6})$$

$$r_j = \frac{1 - \rho}{\rho} \frac{w_j L_j}{K_j} \quad (\text{E.7})$$

$$K_j = \delta \left(1 - \xi + \frac{r_j}{P_j} \right) K_j. \quad (\text{E.8})$$

$$w_j L_j = \sum_{i=0}^J \frac{(\nu_j \tau_{ij})^{-\theta} (\bar{A}_j L_j^\alpha)^\theta}{\sum_{k=1}^J (\nu_k \tau_{ik})^{-\theta} (A_k L_k^\alpha)^\theta} w_i L_i \quad (\text{E.9})$$

$$P_j = \Gamma(1 + (1 - \eta/\theta)) \left(\sum_{i=1}^J (\nu_i \tau_{ji})^{-\theta} (\bar{A}_i L_i^\alpha)^\theta \right)^{-1/\theta} \quad (\text{E.10})$$

$$V_j = \log \left(\bar{\phi}_j L_j^\beta \frac{w_j}{P_j} \right) + \kappa \log \left(\sum_{i=1}^J \exp(\delta V_i - \lambda_{ij})^{1/\kappa} \right) \quad (\text{E.11})$$

$$\mu_{ij} = \frac{\exp(\delta V_i - \lambda_{ij})^{1/\kappa}}{\sum_{i'=1}^J \exp(\delta V_{i'} - \lambda_{i'j})^{1/\kappa}} \quad (\text{E.12})$$

$$L_i = \sum_{j=1}^J \mu_{ij} L_j \quad (\text{E.13})$$

We start with an initial guess of $\{V_j, K_j, w_j\}$ to solve the steady state and then iterate on the system. In particular, at any step of the iteration, we update the variables as follows:

1. Compute μ_{ij} using equation (E.12).
2. Solve L_i as the eigenvector associated with the unit eigenvalue of the matrix $[\mu_{ij}]$, using equation (E.13).
3. Compute r_j using equation (E.7).
4. Compute ν_j using equation (E.5), and P_j using equation (E.10).
5. Solve and update the guess of w_j as a fixed point to equation (E.9).
6. Update V_j using equation (E.11).
7. Update K_j using equation (E.8).

Transition Path To solve a transition path from an initial equilibrium, we use the following iterative algorithm. Start the algorithm with a guess of matrices $\{V_{jt}, K_{jt}\}$. At any step of the iteration, we update the guess as follows:

1. Compute μ_{ijt} forward from $t = 1, 2, \dots, T'$.
2. Starting from period 0, iterative forward with μ_{ijt} to compute the sequence of $\{L_{jt}\}$.
3. Based on $\{L_{jt}, K_{jt}\}$, solve the temporary equilibrium to back out $\{w_{jt}, r_{jt}\}$.
4. Use Equation (E.1) and the initial capital stock to update the guess of $\{K_{jt}\}$ forward from $t = 1, 2, \dots, T'$.
5. Based on the real wage, update V_{jt} backwards from $t = T', T' - 1, \dots, 1$.

E.3 Quantification

We set $\rho = 0.58$, which is the average share of labor compensation in GDP in China between 1995 and 2017 from the Penn World Table. We set the depreciation rate, ξ to be 0.1.

We approximate the initial capital stock in each prefecture in 1994 using the perpetual inventory method. In particular, based on the level of investment in prefecture j in 1994, denoted as I_{j0} , the initial capital stock is computed as $K_{j0} = I_{j0}/\xi$. We measure prefecture-level investment, I_{jt} , using the gross fixed capital formation observed from the *China City Statistical Yearbook*. As the yearbook only started publishing investment flows in 1994, we cannot track the investment series to an earlier date. Moreover, many prefectures did not compile their investment data in the early 1990s. As a result, we only have 216 out of the original 291 prefectures in our sample.