High-Speed Rail and the Gender Employment Gap: Evidence from South Korea*

Eunjee Kwon† Narae Lee‡

March 26, 2026

Abstract

Can place-based infrastructure narrow the gender employment gap? We study South Korea's High-Speed Rail (KTX), a non-commuter intercity service that connected peripheral districts to Seoul between 2004 and 2015. Using a staggered difference-in-differences design, we find that KTX access raised the female-to-male employment ratio in connected non-Seoul districts by 2.8 percentage points (4.2%), with no detectable effect on total employment or population. The mechanism is local sectoral restructuring interacting with a gender-segmented labor market: KTX-connected districts shift from manufacturing toward retail, education, and banking—locally rooted service sectors in which women, who disproportionately re-enter the labor force through sales and service occupations after marriage and childbirth. The effect is spatially concentrated within 10–20 minutes of KTX stations and absent in Seoul, where local service economies were already mature. Our results show that gender-neutral infrastructure can produce gender-differentiated outcomes by transforming the sectoral composition of local labor markets.

*Corresponding author: Eunjee Kwon (kwonee@ucmail.uc.edu). This paper was originally distributed as "Why do improvements in transportation infrastructure reduce the gender gap in South Korea?" Eunjee Kwon is grateful to Matthew E. Kahn, Richard K. Green, and Paulina Oliva for their guidance and encouragement. This paper has benefited substantially from Gabriel M. Ahlfeldt, Vittorio Bassi, Jan Brueckner, Juan Pablo Chauvin, John S. Earle, Matthew Freedman, Miquel-Àngel Garcia López, Edward Glaeser, Marcel Henkel, Naoru Koizumi, Maurice D. Kugler, Jeffery Lin, Pierre Magontier, Andrii Parkomenko, Tanner Regan, Jorge de la Roca, Sita N. Slavov, Daniela Vidart, Elisabet Viladecans-Marsal, and Siqi Zheng, as well as numerous seminar and conference participants. A research grant from the Lusk Real Estate Center and the University of Cincinnati Real Estate Center financially supported this paper. All errors are our own.

†Department of Finance, Lindner College of Business, University of Cincinnati, kwonee@ucmail.uc.edu

‡Development Impact Evaluation (DIME), World Bank Group (IBRD), nlee3@worldbank.org

1 Introduction

What closes the gender employment gap? The dominant explanations center on supply-side forces: expansions of education (Goldin, 2006), reductions in the domestic burden through electrification and appliances (Greenwood, Seshadri, and Yorukoglu, 2005; Greenwood et al., 2016; Vidart, 2024), access to childcare (Maurer-Fazio et al., 2011), and the rise of flexible or remote work arrangements (Goldin, 2014; Harrington and Kahn, 2023). A smaller but growing literature examines demand-side channels, showing that technology shocks to the sectoral composition of labor demand—such as robotization or the shift from brawn to brain—can widen or narrow gender gaps depending on which industries expand (Aksoy, Özcan, and Philipp, 2021; Ngai and Pissarides, 2008; Rendall, 2017). Yet virtually all demand-side evidence comes from economy-wide technological transitions. Whether place-based shocks to the local sectoral mix—the kind that spatial and urban economists routinely study—can reshape gender outcomes remains an open question.

We provide evidence that they can. We study South Korea's High-Speed Rail (KTX), an intercity passenger system that progressively connected Seoul to non-Seoul districts beginning in 2004. Crucially, KTX is not a commuting tool: fewer than 0.1% of passengers use it for daily commuting, and the average one-way trip spans roughly 250 km. Over 80% of trips serve non-commuting purposes—shopping, medical visits, tourism, family visits, and business travel. KTX therefore reduces the cost of moving people but not goods, generating visitor inflows and consumer demand that expand people-facing service sectors—retail, education, banking—while offering no direct stimulus to goods-producing industries such as manufacturing and construction. In South Korea's sharply gender-segmented labor market, where women who re-enter the workforce after marriage and childbirth concentrate overwhelmingly in sales and service occupations (Blau and Kahn, 2017; Chauvin, 2017), this sectoral reallocation amounts to an unintentional demand-side shock to female employment. We find that KTX access raises the female-to-male employment ratio by 2.8 percentage points (4.2%), with no detectable change in total employment, population, or the gender wage gap. The entire effect operates through the between-sector composition of local labor demand: KTX does not make existing employers hire more women, nor does it draw women

into the district—it changes the sectoral mix of jobs in ways that happen to match women's revealed occupational trajectories.

Our main empirical strategy employs the Callaway and Sant'Anna (2021) estimator with staggered treatment timing, defining a district as treated when its centroid falls within 30 minutes' driving time of a KTX station. We restrict attention to non-Seoul districts whose centroids lie within 80 km of a pre-existing (Japanese colonial era) railroad station, ensuring that treated and control districts share similar historical development trajectories. The comparison group includes both not-yet-treated and never-treated districts. Our sample comprises 126 non-Seoul districts (72 treated, 54 never-treated) observed annually from 2000 to 2015.

KTX access increases the female-to-male employment ratio by 2.8 percentage points (4.2%), statistically significant at the 1% level. Total employment and total population are unaffected, establishing that the result reflects a compositional shift—not an aggregate boom. The effect is spatially concentrated: it peaks in the 10–20 minute driving-time ring around KTX stations and dissipates with distance, ruling out diffuse regional trends and supporting a causal interpretation tied to station proximity.

The mechanism operates through sectoral reallocation interacting with gender segmentation. KTX brings increased foot traffic from business travelers, tourists, and day-trippers, while simultaneously granting local residents access to Seoul's amenities—shopping, medical services, and cultural facilities—without permanent relocation. The resulting increase in both visitor flows and residential demand stimulates local service establishments: retail, education, and banking expand significantly, while male-intensive goods-producing sectors—manufacturing, construction, and arts & recreation—contract as the local economy restructures from production toward people-facing services.

This sectoral shift has unintended gender consequences. Because South Korea's labor market is sharply gender-segmented—women who re-enter the workforce after marriage and childbirth disproportionately transition into sales and local service occupations—the service expansion creates job opportunities aligned with the occupational trajectories of post-marriage women. The gender wage gap is unaffected, consistent with a simultaneous expan-

sion of both labor demand (new service-sector jobs) and labor supply (women drawn into the labor force by improved local conditions). Seoul shows no gender employment effect, confirming that the mechanism requires an economy with room for structural transformation from manufacturing toward services.

Results are stable across alternative treatment thresholds (20–60 minutes), control group buffer distances (50–100 km), conditional parallel trends specifications, never-treated-only comparisons, and the Borusyak, Jaravel, and Spiess (2024) imputation estimator. A TWFE specification yields attenuated estimates, consistent with well-known heterogeneity bias in staggered settings (Baker, Larcker, and Wang, 2022).

This paper makes two contributions. First, we show that transportation infrastructure can reshape gender-specific labor market outcomes through the sectoral composition of local labor demand—a channel distinct from the commuting channel that dominates existing work on transportation and gender. Prior studies document how commuting infrastructure widens or narrows gender gaps by differentially expanding men's and women's labor supply: Bütikofer, Løken, and Willén (2024) show that a bridge in Norway widened gender wage gaps by disproportionately expanding men's commuting radii, and Farré, Jofre-Monseny, and Torrecillas (2020), Le Barbanchon, Rathelot, and Roulet (2021), and Gu et al. (2024) document how intra-city commuting costs differentially constrain women. Because KTX is a non-commuter service, this supply-side commuting channel is shut down. Instead, KTX operates through an amenity and consumer-demand channel: it grants non-Seoul residents access to Seoul's amenities without relocation and simultaneously draws visitors into station-adjacent districts, stimulating local service demand. In a gender-segmented labor market, this demand-side sectoral restructuring is inherently gendered—expanding female-intensive services while contracting male-intensive manufacturing—even though the infrastructure itself is gender-neutral.

Second, we provide new evidence on the labor market consequences of high-speed rail, connecting the HSR literature (Lin, 2017; Dong, Zheng, and Kahn, 2020; Qin, 2017) to the urban gender literature (Black, Kolesnikova, and Taylor, 2014; Rosenthal and Strange, 2012; Liu and Su, 2024; Kawabata and Abe, 2018). Our findings also speak to the amenity-based

view of urban growth (Glaeser, Kolko, and Saiz, 2001): KTX redistributes amenity access across space, and the resulting local restructuring has demographic consequences that the amenity literature has not explored. More broadly, we offer a demand-side complement to the supply-side policies—childcare, flexible work—that dominate gender policy discussions (Goldin, 2014; Harrington and Kahn, 2023), showing that place-based infrastructure investments can narrow gender gaps as an unintended byproduct of reshaping local economies.

The remainder of the paper proceeds as follows. Section 2 provides institutional background on South Korea's gender-segmented labor market and the KTX system. Section 3 describes the data. Section 4 details the empirical strategy. Section 5 presents the results, tracing the aggregate and gender employment effects through sectoral reallocation, spatial decay, amenities, migration, and wages. Section 6 documents robustness checks. Section 7 concludes.

2 Background

2.1 Gender-Segmented Labor Market in South Korea

South Korea exhibits one of the largest gender gaps in the labor market among OECD countries (Figure A1). Two features of this gap are central to our analysis: a sharp lifecycle divergence in labor force participation between men and women, and a resulting industry-level gender segmentation that channels men and women into distinct sectors.

Lifecycle Divergence. Men and women in South Korea enter the labor market at comparable rates in their early twenties. A sharp divergence emerges in the mid-twenties, when women begin to marry and have children: many exit the labor force entirely, and those who re-enter do so in their mid-thirties at substantially lower rates (Figure A2). Time-use data from the Korea Labor and Income Panel Study confirm that 40–50% of women report domestic work as their primary activity, compared with fewer than 5% of men (Figure A2).[1]

[1]Despite a shrinking education gap between men and women—female college enrollment surpassed male enrollment in the 2000s—the labor force participation gap has narrowed only modestly, reflecting the persistent role of marriage and childcare in constraining women's employment.

The critical consequence is not merely that fewer women work, but where returning women work: upon re-entry, women disproportionately transition into sales, retail, and local service occupations that offer lower skill-entry barriers, shorter commutes, and flexible work arrangements (Figure A3).

Industry-Level Gender Segmentation. This lifecycle pattern produces a sharply segmented industry structure. Table A2 documents the baseline (2003) female-to-male employment ratio across sectors. Male-intensive sectors include transportation (F/M ratio: 0.13), construction (0.13), and manufacturing (0.62). Female-intensive sectors include restaurants and bars (2.59), hair salons and spas (2.41), medical services (2.06), lodging (1.56), and retail (1.45).[2] The segmentation is structural rather than cyclical: it persists throughout the 2000–2015 sample period and reflects the occupation-switching pattern described above.

This gender segmentation is the key link between infrastructure and labor market outcomes in our analysis. If a transportation investment reshapes the sectoral composition of a local economy—expanding some industries while contracting others—then the gender composition of those industries determines whether the investment disproportionately benefits men or women. We now describe the investment that generates this sectoral shift.

2.2 High-Speed Rail and Non-Seoul Districts

South Korea's 51 million people are heavily concentrated: half reside in the Seoul metropolitan area (Seoul, Incheon, and Gyeonggi provinces), while the remaining half are spread across non-Seoul districts that depend more heavily on manufacturing and agriculture.[3] Non-Seoul districts typically have thinner service-sector ecosystems, fewer female-friendly job opportunities, and weaker connectivity—conditions that motivate female outmigration to Seoul.

[2]The female-to-male ratios are computed from the Census on Establishments at the national level for non-Seoul districts. See Table A2 for detailed sector definitions and KSIC codes.

[3]South Korea's urbanization rate reached 81.7% in 2018, comparable to the United States and France. The Seoul metropolitan area (in Korean, "Sudo-kwon") is defined as the three provinces of Seoul, Incheon, and Gyeonggi (red-bordered area in Figure A4). We use this definition of "Seoul" throughout.

KTX Expansion. The Korea Train eXpress (KTX), introduced in 2004, was designed to reduce regional economic disparities and alleviate congestion pressures in the Seoul area. The system expanded in four major phases (Figure A4; Table A1): 21 stations along the Gyeongbu and Honam lines in April 2004, 9 additional stations on the Gyeongbu and Gyeongui lines in late 2010, 7 stations on the Jeonla line in October 2011, and further extensions in 2012. By 2012, the KTX network covered 22% of the national territory and served 56% of the population.

KTX as a People-Movement Technology. KTX has a defining economic feature: it reduces the cost of moving people across cities while leaving the transportation cost of goods unaffected. The Seoul–Busan trip (325 km) takes six hours by car but only two hours by KTX, with the average one-way KTX trip spanning roughly 250 km. Crucially, KTX is not a commuting tool. Fewer than 0.1% of passengers use it for daily commuting; over 80% of trips serve non-commuting purposes—shopping, education, medical visits, family visits, tourism, and business travel.[4]

For non-Seoul districts, the implications are direct. KTX brings visitors and consumers into local economies—business travelers, tourists, and residents of distant cities accessing local services—generating demand for people-facing sectors such as retail, food and beverage, banking, and education. Because these are precisely the female-intensive sectors identified above (Table A2), KTX creates a Bartik-style sectoral demand shock that is inherently gender-skewed. This channel—intercity connectivity → local service demand → female-intensive sector expansion—is distinct from the commuting channel emphasized in the intra-city transportation literature (Le Barbanchon, Rathelot, and Roulet, 2021; Farré, Jofre-Monseny, and Torrecillas, 2020; Gu et al., 2024). Women in South Korea have shorter commuting distances and are more sensitive to commuting costs than men, but KTX does not expand their commuting radius. Instead, it restructures the local economy in which they already live and work.

[4]A one-way Seoul–Busan KTX ticket costs approximately $40 (2015 prices). With average household disposable income per capita of USD 19,372, the ticket price makes daily commuting prohibitively expensive. The average trip distance of 250 km further confirms that KTX serves intercity, not commuting, demand.

3 Data

We construct a geo-coded balanced panel dataset at the *si-gun-gu* (district) level spanning 16 years (2000–2015).[5] South Korea has 228 districts; after removing islands inaccessible by ground transportation, we retain 218. Restricting to non-Seoul districts whose centroids lie within 80 km of a historical (pre-1920) railroad station yields our preferred analysis sample of 126 districts (72 treated across four KTX cohorts; 54 never-treated controls). We integrate the following data sources:

Census on Establishments. The Census on Establishments is an annual survey conducted by Statistics Korea that covers the universe of approximately 3.3 million business establishments nationwide. Each establishment record contains geographic location (district code), five-digit Korean Standard Industrial Classification (KSIC-9) codes, and total employment counts separately by gender. We aggregate establishment-level data to the district-year-sector level and construct our primary outcome: the female-to-male employment ratio ($\text{emp_f_m}_{it} = \text{emp_fe}_{it}/\text{emp_m}_{it}$), where emp_fe_{it} and emp_m_{it} denote total female and male employment in district i and year t. Secondary outcomes include log female employment ($\log(\text{emp_fe}_{it})$), log male employment ($\log(\text{emp_m}_{it})$), and the female employment share ($\text{share_fe}_{it} = \text{emp_fe}_{it}/\text{emp}_{it}$), which provides a scale-free alternative to the ratio. Sectoral employment shares are constructed by aggregating KSIC-9 codes into 10 broad sectors (see Table A2 for sector definitions and baseline gender composition). Approximately five percent of establishment records have missing worker counts; we impute these from the mean of adjacent years.

Resident Registration Population Survey. The Resident Registration Population Survey is an annual administrative dataset compiled by the Ministry of the Interior and Safety from resident registration records. It provides population counts by gender and age at the

[5] *Si-gun-gu* are South Korea's second-tier administrative divisions, comparable to U.S. counties. The term encompasses three types: *si* (city) and *gun* (rural county), which are subdivisions of provinces; and *gu* (urban district), which are subdivisions of metropolitan cities and Seoul. Numerous boundary changes—mergers, splits, and reclassifications—occurred during our sample period. We harmonize all units to the 2010 administrative boundaries using a crosswalk published by the Seoul Institute (see Appendix B).

district level for every year in our sample. We construct the female-to-male population ratio (pop_f_m_{it}) and log total population.

Korean Labor and Income Panel Study (KLIPS). KLIPS is a longitudinal household survey administered by the Korea Labor Institute. The initial 1998 wave sampled 5,000 households from urban areas; a supplementary sample of 1,415 households from rural areas was added in 2009 to improve geographic representativeness. The survey provides individual-level data on employment, wages, hours worked, and demographic characteristics. We construct district-level weighted average log real monthly wages by gender using both cross-sectional and longitudinal survey weights, deflated to constant 2010 won using the consumer price index (The World Bank, 2024). The relatively small sample (approximately 5,000 wage observations per year with complete weight information) limits statistical power for wage analysis, but KLIPS is the only nationally representative panel wage dataset available for South Korea.

Transportation Data. Driving times from each district centroid to the nearest KTX station are calculated using GIS shortest-path network analysis based on the national road network from the Korea Transport DataBase (KTDB). The OD (origin–destination) travel time matrices are reconstructed for each year to reflect the evolving KTX station network. We also use the locations of railroad stations built during the Japanese colonial era (pre-1920), obtained from historical maps and the Korea Railroad Research Institute, to construct the control group buffer distance measures (see Section 4).

4 Empirical Strategy

4.1 Staggered Difference-in-Differences

We estimate the causal effect of KTX access on the gender employment gap using a difference-in-differences design that exploits the staggered rollout of KTX stations across districts. The

baseline specification takes the form:

$$Y_{it} = \alpha_i + \delta_t + \beta \cdot D_{it} + \varepsilon_{it}, \tag{1}$$

where Y_{it} is the outcome of interest for district i in year t, α_i are district fixed effects, δ_t are year fixed effects, and $D_{it} = \mathbf{1}[t \geq G_i]$ is a binary treatment indicator equal to one once district i's centroid falls within 30 minutes' driving time of an open KTX station. G_i denotes the first year of treatment and ε_{it} is the error term clustered at the district level. Under constant treatment effects, β identifies the average treatment effect on the treated (ATT).

Our primary outcome is the female-to-male employment ratio:

$$Y_{it} = \frac{\text{emp_fe}_{it}}{\text{emp_m}_{it}}, \tag{2}$$

where emp_fe_{it} and emp_m_{it} are total female and male employment in district i at year t from the Census on Establishments. We also estimate effects on the female-to-male population ratio, log female employment, log male employment, and the female employment share ($\text{share_fe}_{it} = \text{emp_fe}_{it}/\text{emp}_{it}$).

Treatment Definition A district i is classified as treated in year t if its geographic centroid falls within 30 minutes' driving time of any KTX station that had opened by year t. The treatment timing variable G_i records the first year a district enters treatment; for never-treated districts, $G_i = \infty$. Four entry cohorts exist: 2004 ($n = 32$), 2010 ($n = 19$), 2011 ($n = 8$), and 2012 ($n = 7$). The original treatment variable identifies an additional cohort entering in 2015 ($n = 6$), the final year of our panel. Because these districts have no post-treatment observations ($t > G_i$), their group-time ATT estimates cannot be identified. Following standard practice in the staggered DiD literature, we reclassify the 2015 cohort as never-treated; in our preferred restricted sample only one district is affected by this reclassification, and results are virtually identical when the 2015 cohort is included. Our preferred control group restriction requires districts to have centroids within 80 km of a pre-1920 railroad station ($\text{km_old_80} = 1$), ensuring geographic comparability. Appendix Figure

A6 maps the treated districts and the old railroad buffer zones.

Preferred Estimator However, recent literature demonstrates that with staggered treatment timing, the standard two-way fixed effects (TWFE) estimator in (1) produces biased estimates when treatment effects vary across cohorts or over time, because already-treated units serve as implicit controls (Baker, Larcker, and Wang, 2022; Goodman-Bacon, 2021; De Chaisemartin and d'Haultfoeuille, 2020). With four treatment cohorts entering between 2004 and 2012, treatment effect heterogeneity is plausible.

To address this, our main specification employs the Callaway and Sant'Anna (2021) estimator, which identifies group-time average treatment effects:

$$ATT(g,t) = E[Y_t(g) - Y_t(0) \mid G = g], \tag{3}$$

where $Y_t(g)$ denotes the potential outcome at time t if first treated in year g, $Y_t(0)$ is the never-treated potential outcome, and G is the cohort indicator. Each $ATT(g,t)$ is estimated from a clean 2×2 DiD comparison, avoiding the contamination that arises in TWFE. We aggregate these into an overall ATT:

$$\widehat{ATT} = \sum_{g} \sum_{t \geq g} \hat{w}(g,t) \cdot \widehat{ATT}(g,t), \tag{4}$$

where the weights $\hat{w}(g,t)$ are proportional to the group size. We also report event-study aggregations, $ATT(e) = \sum_g \hat{w}_g \cdot \widehat{ATT}(g, g+e)$, which plot treatment effects as a function of event-time e (years since initial KTX access), providing a visual test of the parallel trends assumption and the dynamics of the treatment effect.

The main sample comprises 126 non-Seoul districts observed annually from 2000 to 2015. The comparison group includes both not-yet-treated and never-treated districts, and standard errors are clustered at the district level. Pre-trend validity is assessed using the average pre-treatment event-study coefficient and visual inspection of the event-study plot. We report TWFE estimates from (1) as a benchmark in Section 6.

4.2 Control Group Selection

We restrict the sample to non-Seoul districts whose centroids lie within 80 km of a railroad station built during the Japanese colonial era (pre-1920). This restriction ensures that treated and control districts share similar historical development trajectories shaped by over a century of rail proximity, comparable geographic and topological characteristics given Korea's mountainous terrain, and similar strategic considerations in infrastructure placement. This approach follows the logic of using historical transportation networks for identification (Duranton, Morrow, and Turner, 2014; Garcia-López, Holl, and Viladecans-Marsal, 2015; Morten and Oliveira, 2016) while avoiding the stronger exclusion restriction required by an IV design. Since these historical stations predate modern gender labor market dynamics, controlling for their proximity removes confounds related to long-run spatial sorting.

4.3 Alternative Estimators

We complement the main staggered DiD results with several alternative approaches: (i) the Borusyak, Jaravel, and Spiess (2024) imputation estimator; (ii) a standard two-way fixed effects (TWFE) event study specification for comparison; and (iii) sensitivity to treatment thresholds, buffer distances, and control group definitions.

5 Results

5.1 Main Result

Figure 1 and Table 1 present the main results. The event study plots confirm the absence of differential pre-trends and the emergence of treatment effects after KTX entry. Panels (a)–(b) show that log employment and log population are statistically insignificant throughout. KTX access does not generate an aggregate economic boom or contraction in non-Seoul districts. This is consistent with the nature of HSR as a people-movement technology: because KTX reduces the cost of moving people but not goods, it reshapes the sectoral composition of local economies rather than expanding their overall scale. Moreover, any positive demand

stimulus from increased visitor inflows may be offset by a "straw effect," whereby improved Seoul accessibility diverts local consumption toward the capital.

Panels (c)–(d) present the headline gender results. The female-to-male employment ratio rises by 2.8 percentage points (4.2%) after KTX access, statistically significant at the 1% level. The female-to-male population ratio rises by 0.6 percentage points, significant at the 5% level. Together, the null aggregate effects and the significant gender ratio effects establish that KTX induces a compositional shift in local labor markets—not an overall expansion or contraction—redirecting employment toward women relative to men.

What drives this compositional shift? If KTX moves people but not goods, it should expand people-facing service sectors while leaving goods-producing sectors unaffected or worse off. Because South Korea's labor market is sharply gender-segmented—women concentrate in services, men in manufacturing and construction—such sectoral reallocation would mechanically raise the female employment share. The remainder of this section traces this logic through five channels: sectoral employment, spatial decay, amenities, migration, and wages.

5.2 Mechanism

Sectoral Reallocation Table 2 and Figure 2 present the sectoral results. Three male-intensive sectors decline sharply and statistically significantly: manufacturing (−17.3 percent), arts & recreation (−16.3 percent), and construction & utilities (−10.9 percent). Among female-intensive sectors, retail (+9.3 percent), banking & finance (+8.7 percent), and education (+5.4 percent) expand, all statistically significant. Healthcare, restaurants & bars, hair salon & spa, and business & professional services show no statistically significant change.

The pattern is consistent with KTX stimulating people-movement-dependent services at the expense of goods-movement-dependent manufacturing and construction. Because the expanding sectors disproportionately employ women (Table A2), this reallocation mechanically raises the female employment share. Figure 2 visualizes this cross-sectoral gradient: plotting each sector's ATT on log employment against its baseline (2003) female employment share reveals a clear positive gradient. Sectors with higher initial female representation gain

more employment after KTX, while male-dominated sectors contract. Appendix Figure A9 confirms the absence of differential pre-trends in sector-by-sector event study plots.

This sectoral reallocation has a sharp spatial dimension. We exploit variation in driving-time distance to KTX stations to show that the compositional shift is concentrated in station-adjacent areas and dissipates with distance—a pattern that disciplines the causal interpretation by ruling out region-wide trends unrelated to KTX proximity.

Spatial Decay Figure 3 exploits variation in driving-time distance to KTX stations by estimating separate staggered DiD models for each 10-minute driving-time ring. For each ring, districts whose centroids fall within that band are treated; all other districts within the 80 km colonial-railway buffer serve as controls.

Three patterns stand out. First, Panels (a)–(b) show that log employment is statistically insignificant across all rings, while log population rises significantly in the 10–20 minute ring (+5.6 percent). This pattern—more residents but the same number of jobs—suggests that KTX transforms station-adjacent areas from production sites into residential and service hubs.

Second, the gender employment gap effect is spatially concentrated. Panel (c) shows that the female-to-male employment ratio rises by 2.3 percentage points in the 10–20 minute ring, statistically significant, and is indistinguishable from zero elsewhere. The female-to-male population ratio in Panel (d), by contrast, is close to zero in all rings, indicating that the population inflow is gender-balanced; the gender employment gap narrows through labor demand restructuring, not selective female in-migration.

Third, the sectoral channels display clear and complementary spatial gradients. Panel (e) shows that male-intensive sector employment (manufacturing, construction, and arts & recreation) falls sharply in the two rings closest to KTX stations and turns insignificant further out. Panel (f) shows that female-intensive sector employment (retail, banking & finance, and education) peaks in the 10–20 minute ring and turns negative beyond 40 minutes. The 10–20 minute ring is thus the spatial "sweet spot" where male-intensive contraction and female-intensive expansion overlap, producing the strongest gender reallocation. (Appendix Figure A10 reports results for all individual industries.)

The sharp spatial gradient—districts just minutes apart in driving time experience qualitatively different treatment effects—provides additional support for a causal interpretation, as this pattern is difficult to reconcile with unobserved confounders that would need to vary discontinuously at the same spatial frequency as KTX accessibility.

5.3 Other Outcomes

Amenities Following Diamond (2016), we measure amenities as the number of establishments per 10,000 residents in key service sectors. Table 3 reports the results. Retail establishments per capita increase by 3.32 per 10,000 and education facilities by 2.07 per 10,000, both statistically significant. Healthcare facilities show no statistically significant change (-0.01 per 10,000). The retail and education amenity increases are consistent with the sector reallocation results and may have reduced women's labor force participation costs by improving access to childcare and daily services.

Wage Channel Using KLIPS wage data, we estimate the effect on log real wages by gender. Table 4 reports the results. All estimates—female wages, male wages, and the gender wage gap—are statistically insignificant across both current and long-run real wage measures. The null wage effect, combined with the increase in female employment share, is consistent with a simultaneous expansion of both labor demand and labor supply in female-intensive sectors: KTX-induced service growth creates new jobs (demand shift), while the improved local amenities and accessibility attract women into the labor force (supply shift). Because both curves shift outward, employment rises while wages remain stable—ruling out a pure demand story (which would raise wages) or a pure supply story (which would lower them).

5.4 Discussion

The results above point to a unified mechanism. KTX does not simply scale up local economies—it restructures them. Station-adjacent districts (10–20 minutes) experience population growth without employment growth, suggesting that KTX transforms these areas

into more attractive residential locations—in part because residents can now access Seoul's amenities (shopping, medical services, cultural facilities) via a short KTX trip without relocating to the capital. At the same time, the expansion of retail, education, and banking establishments—the amenity improvements documented in Section 5.3—enhances local quality of life, making these districts better places to live even as total job counts remain flat.

The restructuring unintentionally gives more opportunities to women, given the pre-existing labor market structure in South Korea. The KTX-induced growth of retail, education, and financial services in non-Seoul districts—particularly those serving residential demand and benefiting from increased foot traffic—creates job opportunities that align with the occupational trajectories of post-marriage women (Section 2.1), providing a micro-foundation for why an ostensibly gender-neutral infrastructure investment produces gender-differentiated employment effects. The null wage result is consistent with this interpretation: both labor demand (new service-sector jobs) and labor supply (women drawn into the labor force by improved local conditions) expand simultaneously, raising female employment while leaving wages unchanged.

A final piece of evidence comes from Seoul. Running the same specification on Seoul metropolitan districts yields a statistically insignificant employment ratio ATT of -0.009 (Appendix Figure A11). Seoul's economy was already dominated by services before KTX—the sectors that expand in non-Seoul districts (retail, education, banking) were already mature in Seoul, leaving no room for the compositional shift that drives the gender effect elsewhere. This confirms that the mechanism is specific to economies with room for structural transformation from goods-producing to people-facing sectors.

6 Robustness Checks

We subject the main finding—KTX raises the female-to-male employment ratio in non-Seoul districts—to an extensive battery of robustness checks. Table 5 collects the ATT estimates across all specifications; Figures 4–5 provide visual evidence.

Treatment Thresholds. Our baseline defines treatment as a 30-minute driving-time threshold. We re-estimate the model using 20, 45, and 60 minutes. Table 5, Panel A reports the results: all four thresholds produce positive and statistically significant ATTs, ranging from 0.020 (45 minutes) to 0.028 (30 and 60 minutes). The stronger effects at tighter thresholds (20 and 30 minutes) and the modest attenuation at 45 minutes are consistent with proximity to KTX stations mattering, as expected from the spatial decay evidence in Section 5.2. Figure 4, Panel (a) overlays the event study plots for all four thresholds, confirming that the post-treatment divergence is robust to the choice of cutoff.

Control Group Buffers. Our baseline restricts the sample to districts within 80 km of a colonial-era railroad station. We vary this buffer to 50, 70, and 100 km. The ATT is virtually unchanged across all specifications (0.028–0.029; Table 5, Panel B), confirming that the results are not driven by the particular geographic comparison set or by the inclusion of remote districts with different development trajectories.

Conditional Parallel Trends. We re-estimate the model conditioning on time-varying covariates. The ATT is identical to the unconditional specification (0.028; Table 5, Panel C). This is expected in our balanced panel: the Callaway and Sant'Anna (2021) regression-adjusted estimator already absorbs district fixed effects, rendering time-invariant or slowly-moving covariates redundant. Figure 4, Panel (b) overlays the conditional and unconditional event study plots, which are indistinguishable.

Never-Treated Only. Our baseline comparison group includes both not-yet-treated and never-treated districts. Restricting to never-treated districts only (dropping all not-yet-treated units) produces a statistically significant ATT of 2.7 percentage points (Table 5, Panel C), virtually identical to the baseline. This rules out concerns about anticipation effects or contamination from the future treatment status of not-yet-treated units.

Alternative Estimators. We complement the Callaway and Sant'Anna (2021) estimator with two alternatives. First, the Borusyak, Jaravel, and Spiess (2024) imputation estimator

yields a positive average ATT of 0.011, with individual horizon estimates positive and τ_4 significant at 5% (Table 5, Panel C). Second, a standard TWFE event study with relative-time indicators and district and year fixed effects produces pre-treatment coefficients clustered around zero and positive, increasing post-treatment coefficients (Figure 5). The TWFE post-treatment ATT (0.012) is attenuated relative to the staggered DiD estimate, consistent with the well-documented negative weighting bias in staggered settings (Baker, Larcker, and Wang, 2022), but the qualitative pattern is identical.

Alternative Outcome. We replace the female-to-male employment ratio with the female employment share (share_fe = emp_fe/emp), a scale-free measure that removes district-size noise. The ATT on share_fe is +0.9 percentage points, statistically significant at the 1% level (Table 5, Panel D). Male employment share moves symmetrically (−0.9 percentage points, also significant). This confirms that the main result is not an artifact of the ratio functional form.

7 Conclusion

This paper shows that South Korea's High-Speed Rail unintentionally narrowed the gender employment gap in connected non-Seoul districts. The mechanism operates through sectoral reallocation: KTX transforms station-adjacent districts into more attractive residential and commercial locations, stimulating service establishments—retail, education, and banking—while male-intensive manufacturing and construction contract. In South Korea's gender-segmented labor market, this restructuring disproportionately benefits women, who re-enter the workforce through precisely the sales and service occupations that KTX-induced demand creates. The effect operates entirely on the composition margin—total employment and population are unaffected—and is absent in Seoul, where service economies were already mature. The null wage effect is consistent with simultaneous expansion of labor demand and supply in female-intensive sectors.

Our findings carry three broader implications. First, infrastructure investments that move people rather than goods can have unintended gendered consequences through their differ-

ential impact on the sectoral composition of local economies. Policymakers evaluating large transport projects should account for these compositional effects, particularly in economies with gender-segmented labor markets. Second, the mechanism we document—connectivity-induced structural change from goods-producing toward people-facing services—may be relevant for other countries investing in intercity rail or highway systems, where similar gender segmentation prevails. Third, our results complement the commuting-focused gender-transport literature (Le Barbanchon, Rathelot, and Roulet, 2021; Bütikofer, Løken, and Willén, 2024) by identifying a distinct channel: even non-commuter intercity rail can affect gender gaps through local economic restructuring rather than commuting radius expansion.

An important implication is that the narrowing of the gender employment gap does not necessarily imply an improvement in gender equality. The expanding service-sector jobs—retail, personal care, and education—tend to offer lower wages and fewer career advancement opportunities than the male-dominated sectors that contract. Our wage analysis shows no compression of female wages, but neither does it show convergence toward male wage levels. The compositional shift may therefore reflect a deepening of existing occupational segmentation rather than a broadening of women's economic opportunities.

Several limitations suggest directions for future research. Our district-level analysis cannot disentangle individual-level job transitions or identify which women are affected. Census microdata linking residential and workplace locations could provide richer evidence on the individual margin (Venator, 2022; Fan and Zou, 2021; Costa and Kahn, 2000). KTX may also interact with channels we cannot fully isolate—such as access to grandparents for childcare (Anstreicher and Venator, 2024) or changes in couples' joint location decisions (Gu et al., 2024). Finally, a structural model integrating intercity connectivity with gender-heterogeneous labor supply and demand would help quantify welfare effects and guide optimal infrastructure policy design.

References

Aksoy, Cevat Giray, Berkay Özcan, and Julia Philipp (2021). "Robots and the gender pay gap in Europe". In: *European Economic Review* 134, p. 103693.

Anstreicher, Garrett and Joanna Venator (2024). "To Grandmother's House We Go: Informal Childcare and Female Labor Mobility". In.

Baker, Andrew C., David F. Larcker, and Charles C.Y. Wang (2022). "How much should we trust staggered difference-in-differences estimates?" In: *Journal of Financial Economics* 144.2, pp. 370–395.

Black, Dan A, Natalia Kolesnikova, and Lowell J Taylor (2014). "Why do so few women work in New York (and so many in Minneapolis)? Labor supply of married women across US cities". In: *Journal of Urban Economics* 79, pp. 59–71.

Blau, Francine D and Lawrence M Kahn (2017). "The gender wage gap: Extent, trends, and explanations". In: *Journal of Economic Literature* 55.3, pp. 789–865.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess (2024). "Revisiting Event-Study Designs: Robust and Efficient Estimation". In: *Review of Economic Studies* 91.6, pp. 3253–3285.

Bütikofer, Aline, Katrine V. Løken, and Alexander Willén (May 2024). "Building Bridges and Widening Gaps". In: *The Review of Economics and Statistics* 106.3, pp. 681–697.

Callaway, Brantly and Pedro HC Sant'Anna (2021). "Difference-in-differences with multiple time periods". In: *Journal of econometrics* 225.2, pp. 200–230.

Chauvin, Juan Pablo (2017). "Gender-Segmented Labor Markets and the Effects of Local Demand Shocks". In: *Job Market Paper, Harvard University*.

Costa, Dora L and Matthew E Kahn (2000). "Power couples: changes in the locational choice of the college educated, 1940–1990". In: *The Quarterly Journal of Economics* 115.4, pp. 1287–1315.

De Chaisemartin, Clément and Xavier d'Haultfoeuille (2020). "Two-way fixed effects estimators with heterogeneous treatment effects". In: *American economic review* 110.9, pp. 2964–2996.

Diamond, Rebecca (2016). "The determinants and welfare implications of US workers' diverging location choices by skill: 1980-2000". In: *American Economic Review* 106.3, pp. 479–524.

Dong, Xiaofang, Siqi Zheng, and Matthew E Kahn (2020). "The role of transportation speed in facilitating high skilled teamwork across cities". In: *Journal of Urban Economics* 115, p. 103212.

Duranton, Gilles, Peter M Morrow, and Matthew A Turner (2014). "Roads and Trade: Evidence from the US". In: *Review of Economic Studies* 81.2, pp. 681–724.

Fan, Jingting and Ben Zou (2021). "The dual local markets: Family, jobs, and the spatial distribution of skills". In: *SSRN Electronic Journal,(November)*.

Farré, Lídia, Jordi Jofre-Monseny, and Juan Torrecillas (2020). "Commuting Time and the Gender Gap in Labor Market Participation". In.

Garcia-López, Miquel-Àngel, Adelheid Holl, and Elisabet Viladecans-Marsal (2015). "Suburbanization and highways in Spain when the Romans and the Bourbons still shape its cities". In: *Journal of Urban Economics* 85, pp. 52–67.

Glaeser, Edward L, Jed Kolko, and Albert Saiz (2001). "Consumer city". In: *Journal of Economic Geography* 1.1, pp. 27–50.

Goldin, Claudia (2006). "The quiet revolution that transformed women's employment, education, and family". In: *American economic review* 96.2, pp. 1–21.

— (2014). "A grand gender convergence: Its last chapter". In: *American economic review* 104.4, pp. 1091–1119.

Goodman-Bacon, Andrew (2021). "Difference-in-differences with variation in treatment timing". In: *Journal of econometrics* 225.2, pp. 254–277.

Greenwood, Jeremy, Ananth Seshadri, and Mehmet Yorukoglu (2005). "Engines of liberation". In: *The Review of Economic Studies* 72.1, pp. 109–133.

Greenwood, Jeremy et al. (2016). "Technology and the changing family: A unified model of marriage, divorce, educational attainment, and married female labor-force participation". In: *American Economic Journal: Macroeconomics* 8.1, pp. 1–41.

Gu, Yizhen et al. (2024). "Home location choices and the gender commute gap". In: *Journal of Human Resources* 59.2, pp. 545–575.

Harrington, Emma and Matthew E Kahn (2023). *Has the rise of work-from-home reduced the motherhood penalty in the labor market.*

Kawabata, Mizuki and Yukiko Abe (2018). "Intra-metropolitan spatial patterns of female labor force participation and commute times in Tokyo". In: *Regional Science and Urban Economics* 68, pp. 291–303.

Le Barbanchon, Thomas, Roland Rathelot, and Alexandra Roulet (2021). "Gender differences in job search: Trading off commute against wage". In: *The Quarterly Journal of Economics* 136.1, pp. 381–426.

Lin, Yatang (2017). "Travel costs and urban specialization patterns: Evidence from China's high speed railway system". In: *Journal of Urban Economics* 98, pp. 98–123.

Liu, Sitian and Yichen Su (2024). "The geography of jobs and the gender wage gap". In: *Review of Economics and Statistics* 106.3, pp. 872–881.

Maurer-Fazio, Margaret et al. (2011). "Childcare, eldercare, and labor force participation of married women in urban China, 1982–2000". In: *Journal of Human Resources* 46.2, pp. 261–294.

Morten, Melanie and Jaqueline Oliveira (2016). *Paving the way to development: Costly migration and labor market integration.* Tech. rep. National Bureau of Economic Research.

Ngai, L Rachel and Christopher A Pissarides (2008). "Trends in hours and economic growth". In: *Review of Economic dynamics* 11.2, pp. 239–256.

Qin, Yu (2017). "'No county left behind?'The distributional impact of high-speed rail upgrades in China". In: *Journal of Economic Geography* 17.3, pp. 489–520.

Rendall, Michelle (2017). "Brain versus brawn: the realization of women's comparative advantage". In: *University of Zurich, Institute for Empirical Research in Economics, Working Paper* 491.

Rosenthal, Stuart S and William C Strange (2012). "Female entrepreneurship, agglomeration, and a new spatial mismatch". In: *Review of Economics and Statistics* 94.3, pp. 764–788.

The World Bank (2024). *World Development Indicators: Consumer Price Index (2010 = 100)*. Retrieved from https://databank.worldbank.org/source/world-development-indicators.

Venator, Joanna (2022). *Dual-earner migration decisions, earnings, and unemployment insurance*. Boston College.

Vidart, Daniela (2024). "Human capital, female employment, and electricity: Evidence from the early 20th-century United States". In: *Review of Economic Studies* 91.1, pp. 560–594.

Table 1: Impact of KTX on Employment and Population (Non-Seoul)

	(1) F/M Emp Ratio	(2) F/M Pop Ratio	(3) Log Emp	(4) Log Pop
ATT	0.028***	0.006**	0.014	0.025
	(0.008)	(0.003)	(0.026)	(0.020)
Pre-treatment mean	0.671	1.010	9.84	11.08
Districts	126	126	126	126

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Each column reports the Callaway and Sant'Anna (2021) simple ATT aggregation. Treatment: district centroid within 30 minutes' driving time of a KTX station. Comparison group: not-yet-treated and never-treated districts within 80 km of a colonial-era railroad station. Standard errors in parentheses, clustered at the district level. Pre-treatment means are for treated districts in 2003. Sample: 126 non-Seoul districts (72 treated, 54 never-treated), 2000–2015.

Table 2: Sectoral Employment Effects of HSR (Non-Seoul)

Sector	(1) ATT	(2) SE
Panel A: Male-Intensive Sectors		
Construction & Utilities	−0.109***	(0.038)
Business & Professional	0.046	(0.078)
Arts & Recreation	−0.163***	(0.050)
Manufacturing	−0.173***	(0.040)
Panel B: Female-Intensive Sectors		
Banking & Finance	0.087**	(0.036)
Retail	0.093***	(0.029)
Education	0.054**	(0.025)
Healthcare	0.033	(0.039)
Hair Salon & Spa	−0.036	(0.023)
Restaurants & Bars	−0.005	(0.023)

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Each row reports the staggered DiD ATT on the sector's log employment within a district. Sectors are sorted by baseline (2003) female share of total employment within each panel (ascending). Panel A includes sectors where women constitute less than half of the workforce; Panel B includes sectors where women are the majority. Standard errors in parentheses, clustered at the district level. Sample: Non-Seoul districts within 80 km of a historical railroad station. See Table A2 for sector definitions and baseline gender composition.

Table 3: Amenity Channel: Establishments per 10,000 Residents (Non-Seoul)

Outcome	(1) ATT	(2) SE	(3) p-value	(4) Sig.
Retail (General, G471) per 10k pop	3.324	(0.562)	<0.001	***
Education (P) per 10k pop	2.068	(0.891)	0.020	**
Healthcare (Q86) per 10k pop	−0.010	(0.326)	0.977	

Notes: Staggered DiD ATT estimates. Each row reports the effect of KTX access on the number of establishments per 10,000 residents in the indicated sector. Standard errors in parentheses, clustered at the district level. Sample: Non-Seoul districts within 80 km of historical railroad stations.

Table 4: Wage Channel: Impact of HSR on Log Real Wages (KLIPS, Non-Seoul)

Outcome	(1) ATT	(2) SE	(3) p-value	(4) Sig.
F/M wage ratio (current)	0.013	(0.061)	0.830	
F/M wage ratio (long-run)	−0.012	(0.061)	0.842	
Log wage gap, F−M (current)	0.081	(0.081)	0.318	
Log wage gap, F−M (long-run)	0.056	(0.081)	0.491	

Notes: Staggered DiD ATT estimates. "Current" wages are deflated monthly wages in the survey year; "long-run" wages use the permanent income measure from KLIPS. F/M wage ratio = female wage / male wage. Log wage gap = log(female wage) − log(male wage). Sample: Non-Seoul districts within 80 km of historical railroad stations. Panel is unbalanced ($N = 1{,}198$ district-year pairs with non-missing wage data). Standard errors clustered at the district level.

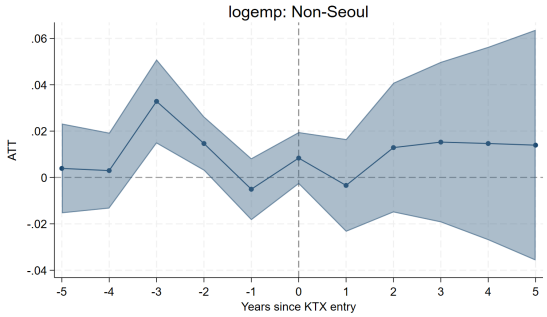
Table 5: Robustness Summary: ATT on Female-to-Male Employment Ratio across Specifications

Specification	(1) ATT	(2) SE	(3) p-value	(4) Sig.
Panel A: Treatment Threshold Sensitivity				
Main (30 min, baseline)	0.028	(0.008)	0.001	***
20 min threshold	0.026	(0.008)	0.001	***
45 min threshold	0.020	(0.009)	0.031	**
60 min threshold	0.028	(0.009)	0.002	***
Panel B: Control Group Buffer Distance				
Main (80 km, baseline)	0.028	(0.008)	0.001	***
50 km buffer	0.028	(0.009)	0.002	***
70 km buffer	0.028	(0.008)	0.001	***
100 km buffer	0.029	(0.008)	<0.001	***
Panel C: Estimator and Specification				
Conditional PT	0.028	(0.008)	0.001	***
Never-treated only	0.027	(0.009)	0.002	***
Borusyak imputation (avg.)	0.011	—	—	†
TWFE post-dummy	0.012	(0.008)	0.113	
Panel D: Alternative Outcome				
Female emp. share (share_fe)	0.009	(0.003)	0.001	***
Male emp. share (share_m)	−0.009	(0.003)	0.001	***

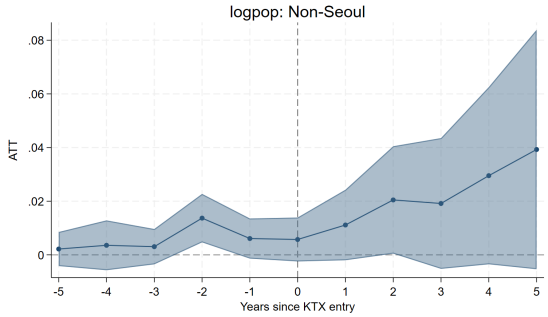
Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Each row reports the overall (simple) staggered DiD ATT unless otherwise noted. All specifications use the non-Seoul sample. Panels A–B vary the treatment threshold and control group buffer while holding other parameters at baseline values. Panel C varies the estimator or comparison group. † Borusyak ATT is the average of horizon-specific τ_0–τ_5 estimates from the Borusyak, Jaravel, and Spiess (2024) imputation estimator; individual horizons are positive, with τ_4 significant at 5%. TWFE uses a simple post-treatment dummy with district and year fixed effects. Standard errors in parentheses, clustered at the district level.

Figure 1: Impact of HSR on Local Outcomes in Non-Seoul, Event Study

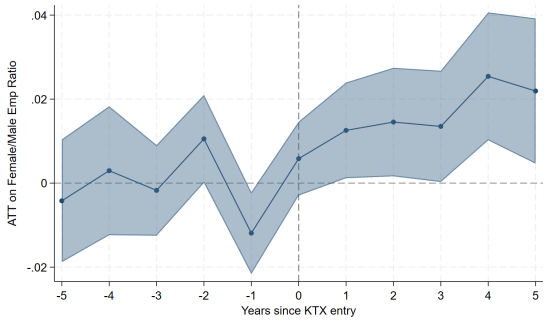
(a) Log Employment



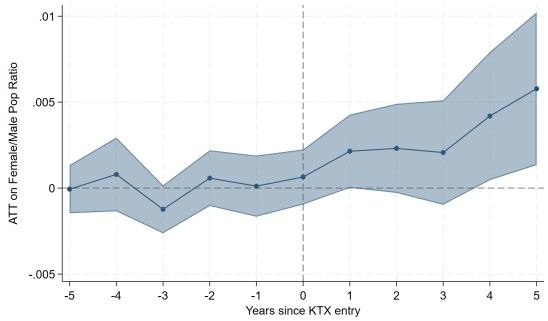
(b) Log Population



(c) Female-to-Male Emp Ratio

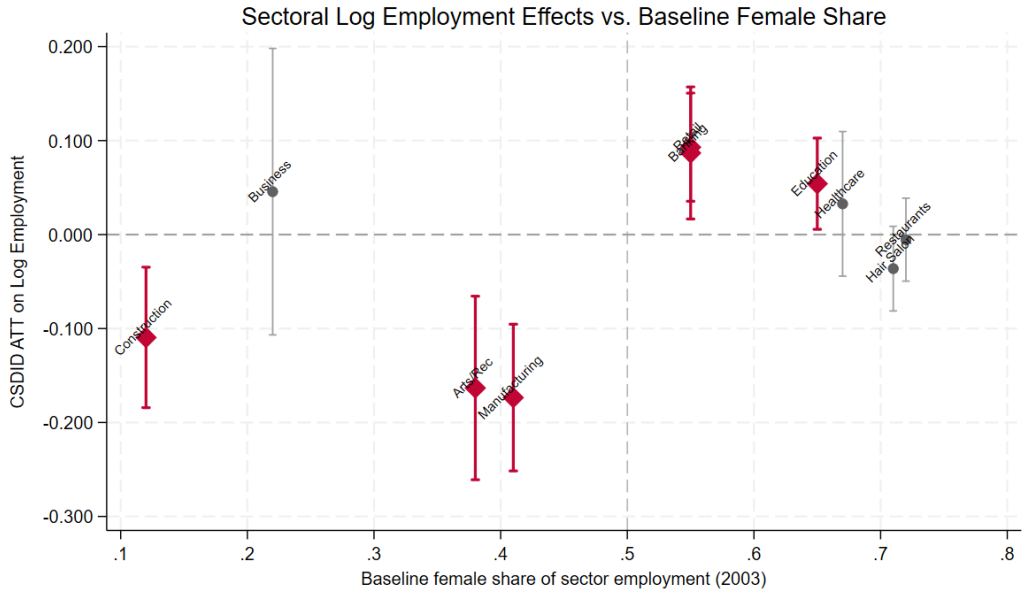


(d) Female-to-Male Pop Ratio



Notes: Each panel plots the staggered difference-in-differences event study estimates of the ATT on the indicated outcome variable. The horizontal axis denotes years relative to initial KTX access; the vertical axis shows the estimated treatment effect. Panels (a)–(b) show aggregate outcomes (log employment and log population); Panels (c)–(d) show gender composition outcomes (female-to-male employment ratio and population ratio). The comparison group includes both not-yet-treated and never-treated districts. Dashed horizontal lines mark zero. 95% confidence intervals shown. Standard errors clustered at the district level. Sample: Non-Seoul districts within 80 km of a historical railroad station ($N = 126$ districts, 2000–2015).

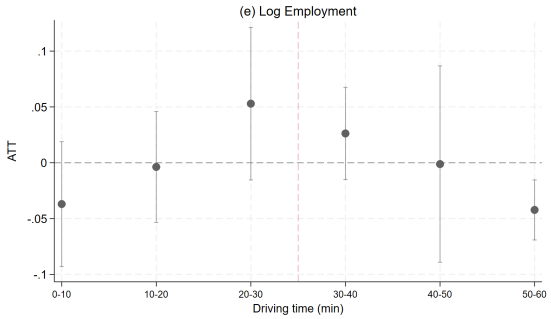
Figure 2: Sectoral Effects of HSR vs. Baseline Female Employment Share



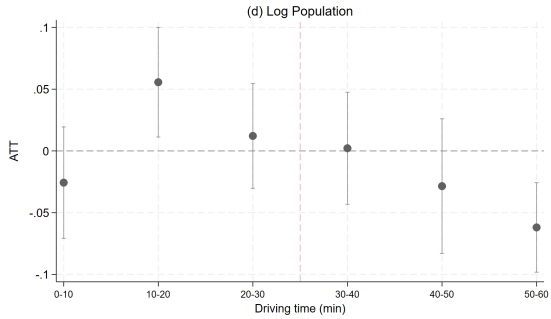
Notes: Each point represents one of ten sectors. The horizontal axis measures the sector's baseline (2003) female share of total employment, computed from the Census of Establishments. The vertical axis plots the staggered DiD ATT on each sector's log employment level. Red diamonds denote statistical significance at the 5% level; gray circles denote insignificant effects. Vertical bars represent 95% confidence intervals. The positive gradient indicates that female-intensive sectors expand (or contract less) after KTX entry, while male-dominated sectors shrink—consistent with KTX-driven demand shifts toward people-facing service industries. Sample: Non-Seoul districts within 80 km of a historical railroad station. See Table 2 for point estimates and Table A2 for sector definitions.

Figure 3: Spatial Decay of KTX Effects by Driving Time

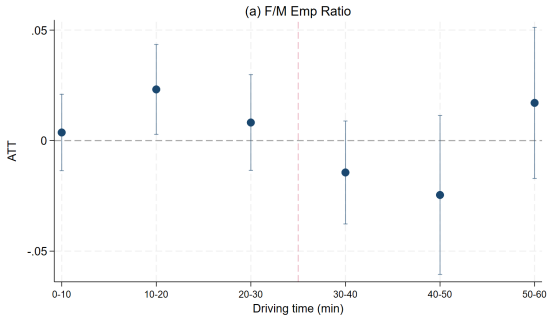
(a) Log Employment



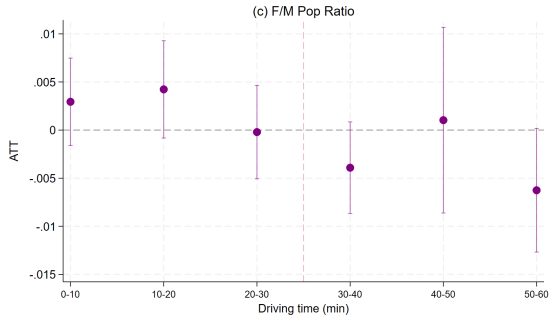
(b) Log Population



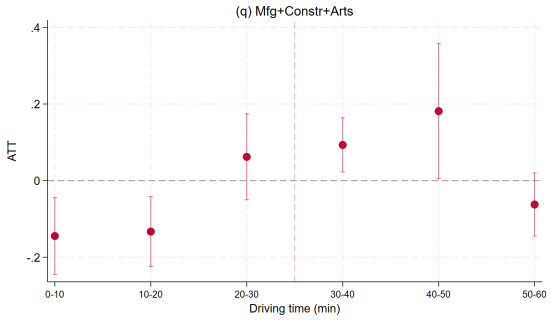
(c) Female-to-Male Emp Ratio



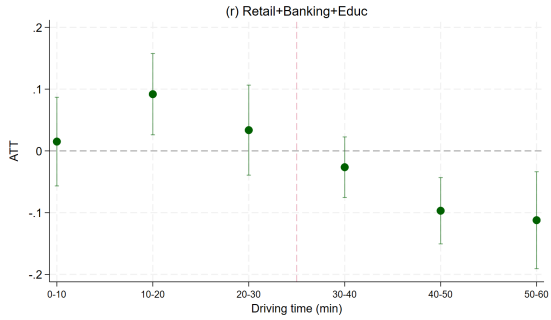
(d) Female-to-Male Pop Ratio



(e) Male-Intensive Sectors



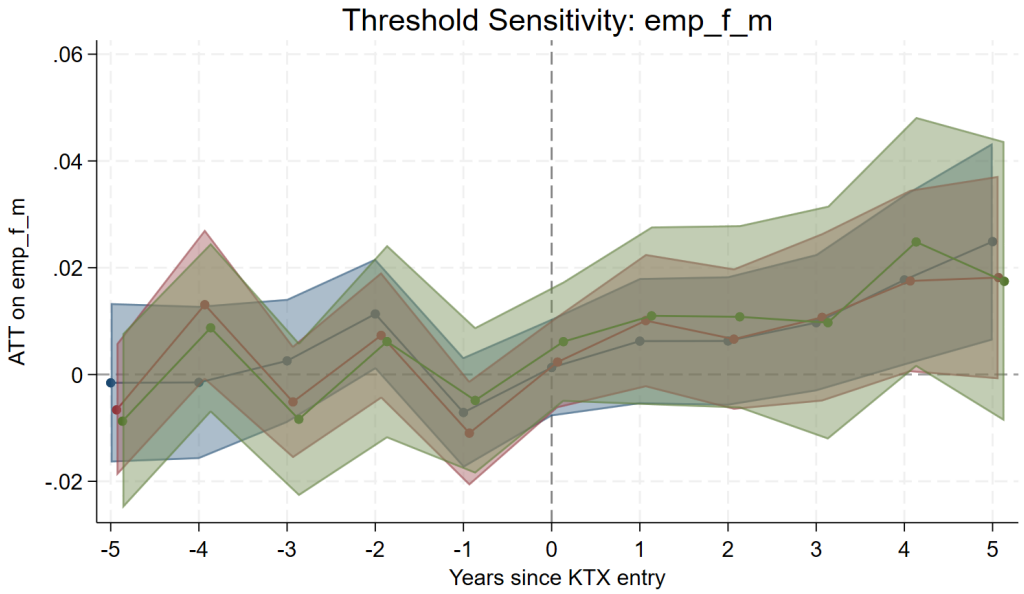
(f) Female-Intensive Sectors



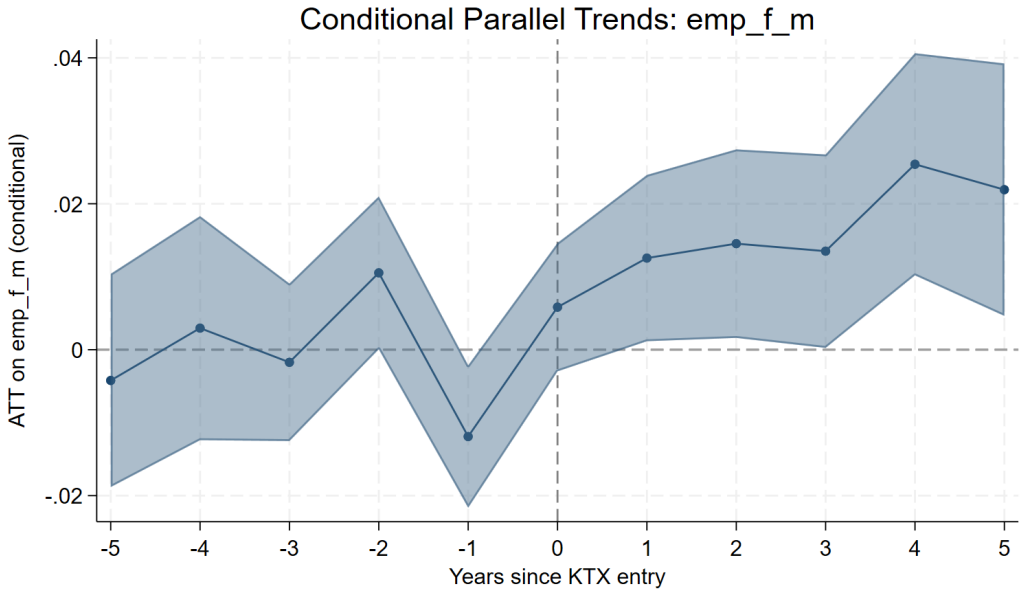
Notes: Each panel plots staggered DiD ATT estimates by driving-time ring (10-minute bins) from the nearest KTX station. The horizontal axis denotes the driving-time ring; the vertical axis shows the estimated treatment effect. For each ring, treated districts are those whose centroids fall within the indicated driving-time band; controls are all other districts within the 80 km colonial-railway buffer (not-yet-treated as comparison). Panels (a)–(b) show aggregate outcomes; Panels (c)–(d) show gender composition outcomes; Panels (e)–(f) show sectoral employment effects. Male-intensive sectors in Panel (e) combine manufacturing, construction & utilities, and arts & recreation—the three sectors with statistically significant negative ATTs in Table 2. Female-intensive sectors in Panel (f) combine retail, banking & finance, and education—the three sectors with statistically significant positive ATTs. See Table A2 for baseline gender composition by sector. Bars represent 95% confidence intervals. Standard errors clustered at the district level. Red dashed line marks the 30-minute main treatment threshold.

Figure 4: Robustness: Threshold Sensitivity and Conditional Parallel Trends

(a) Threshold Sensitivity

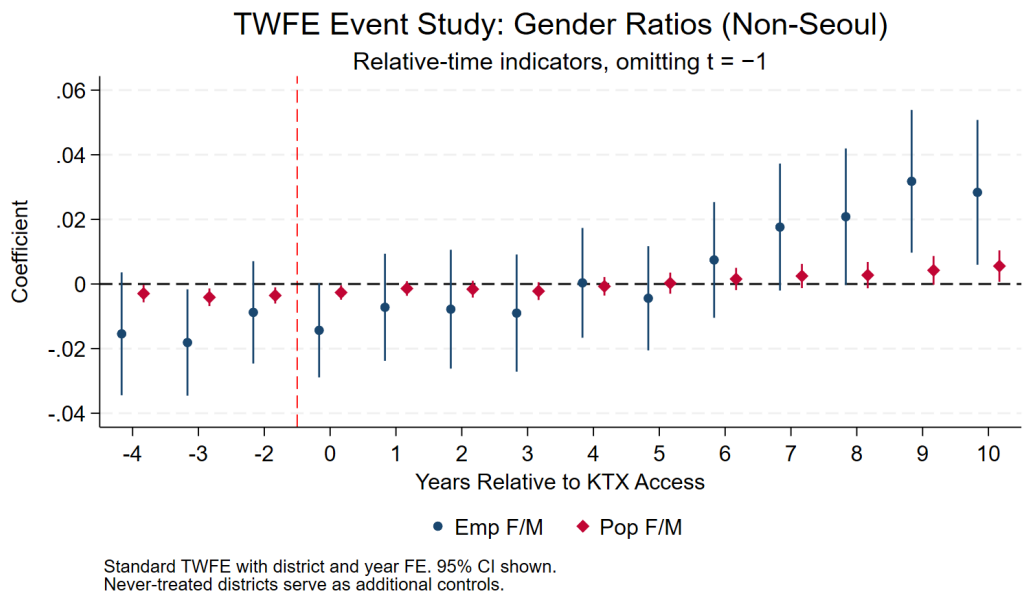


(b) Conditional Parallel Trends



Notes: Panel (a): Staggered DiD event study estimates of the ATT on the female-to-male employment ratio under alternative treatment thresholds (20, 30, 45, 60 minutes driving time to the nearest KTX station). All specifications use the 80 km historical railroad buffer and not-yet-treated comparison group. Panel (b): Comparison of staggered DiD event study estimates with and without conditioning on time-varying covariates (`xformla`). The near-identical estimates confirm that conditional and unconditional parallel trends yield the same result in this balanced panel. 95% confidence intervals shown. Standard errors clustered at the district level.

Figure 5: TWFE Event Study: Gender Ratios (Non-Seoul)

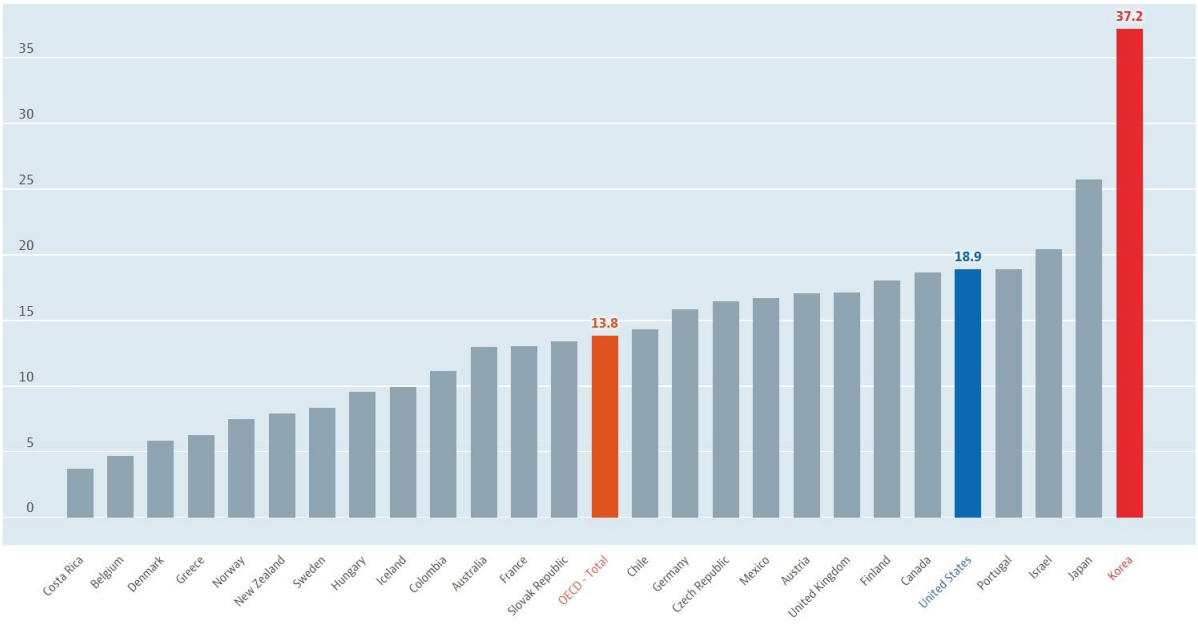


Notes: Standard TWFE event study with district and year fixed effects. Relative-time indicators from $t = -4$ to $t = +10$, omitting $t = -1$ as the reference period. Never-treated districts serve as additional controls (all indicators equal zero). 95% confidence intervals shown. Standard errors clustered at the district level. Sample: Non-Seoul districts within 80 km of historical railroad stations.

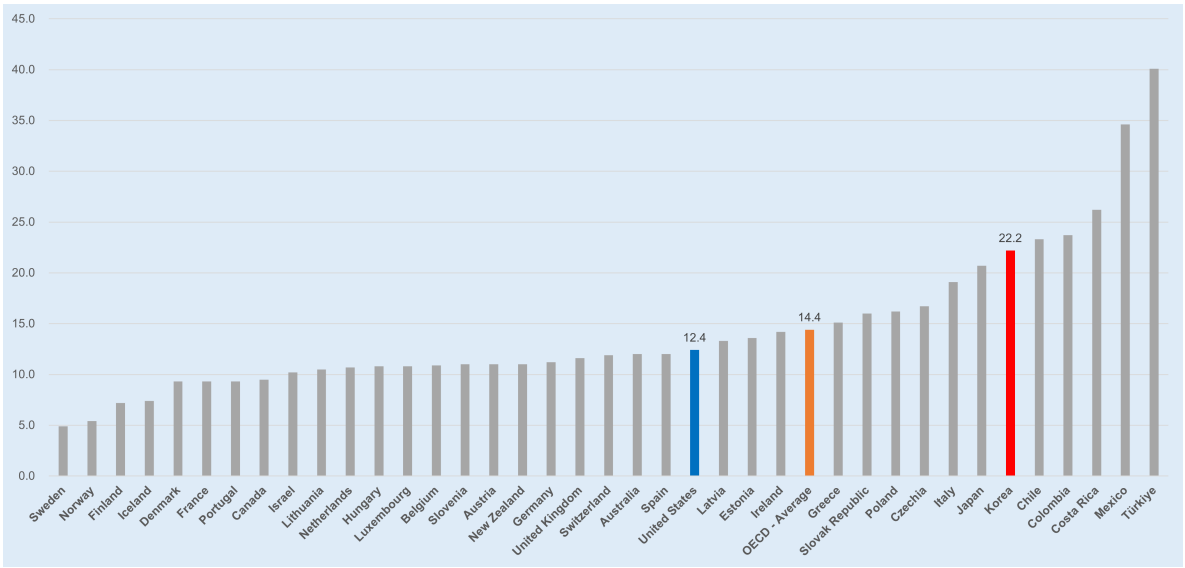
A Background Figures

Figure A1: Cross-Country Gender Gap in the Labor Market

(a) Gender Wage Gap (Percentage Points)



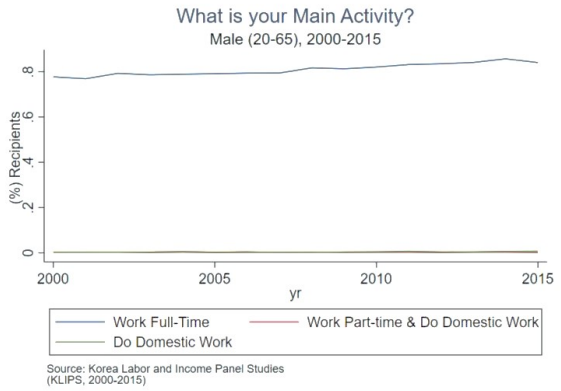
(b) Labor Force Participation Gap (Percentage Points)



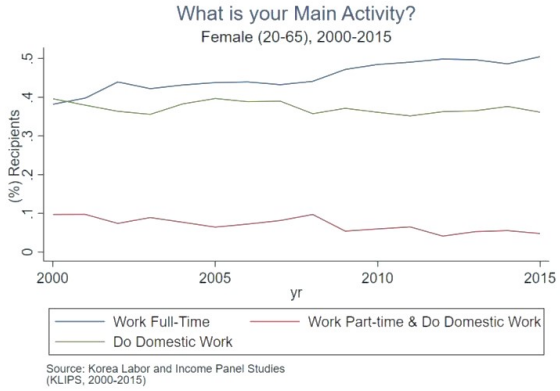
Notes: The gender wage gap in panel (a) is the difference between the median earnings of females and males. The labor force participation gap in panel (b) is the difference in the labor participation rate of males and females. Source: OECD (2020).

Figure A2: Gender Differences in Time Use and Life Cycle

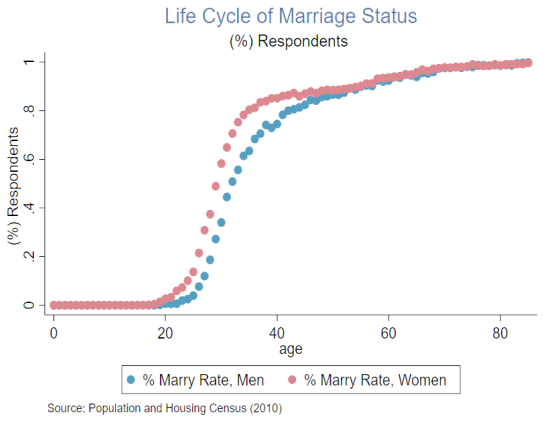
(a) Time Use: Men



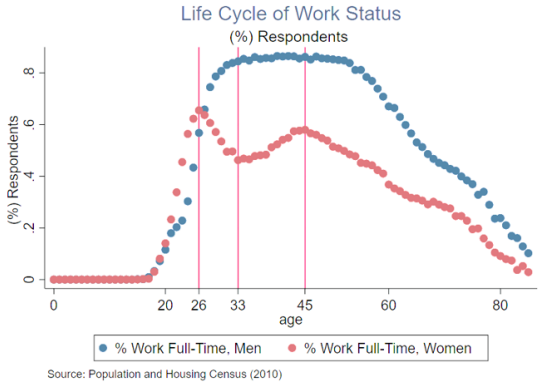
(b) Time Use: Women



(c) Life Cycle: Married (%)



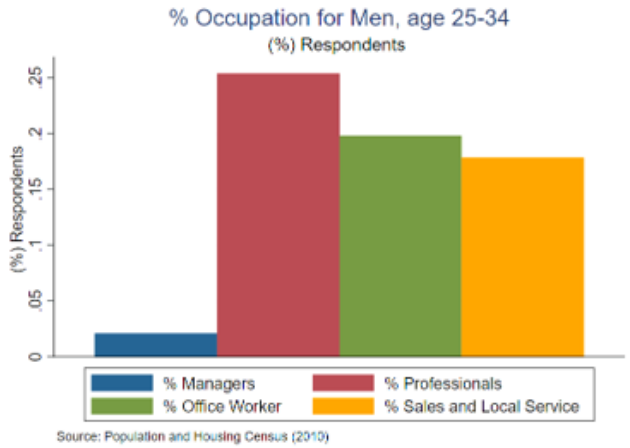
(d) Life Cycle: Labor Force Participation (%)



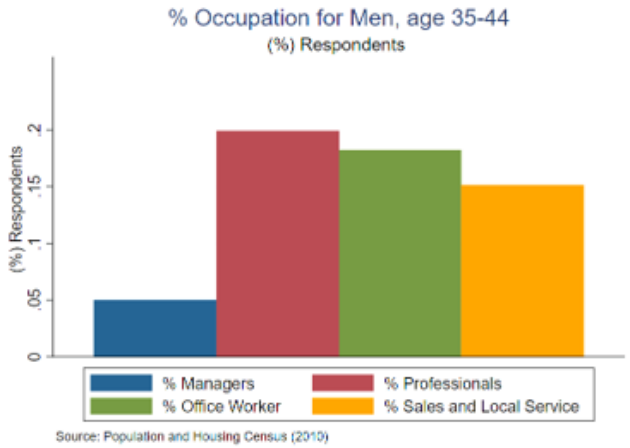
Notes: Panels (a)–(b): Annual share of survey answers to "What is your main activity last month?" from KLIPS, for men and women. Panels (c)–(d): Life cycle of men (blue) and women (red) with respect to marriage and labor force participation, Resident Registration Population Survey (2010).

Figure A3: Occupation Composition for Gender and Age Groups

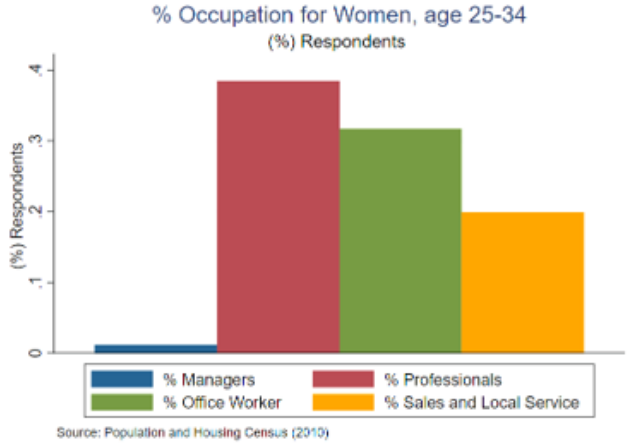
(a) Men (age 25-34)



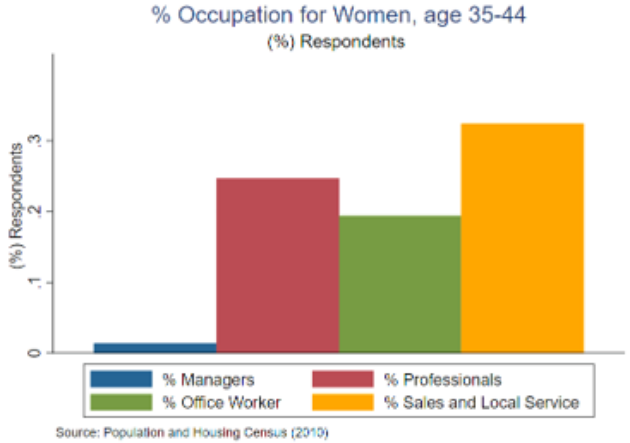
(b) Men (age 35-44)



(c) Women (age 25-34)



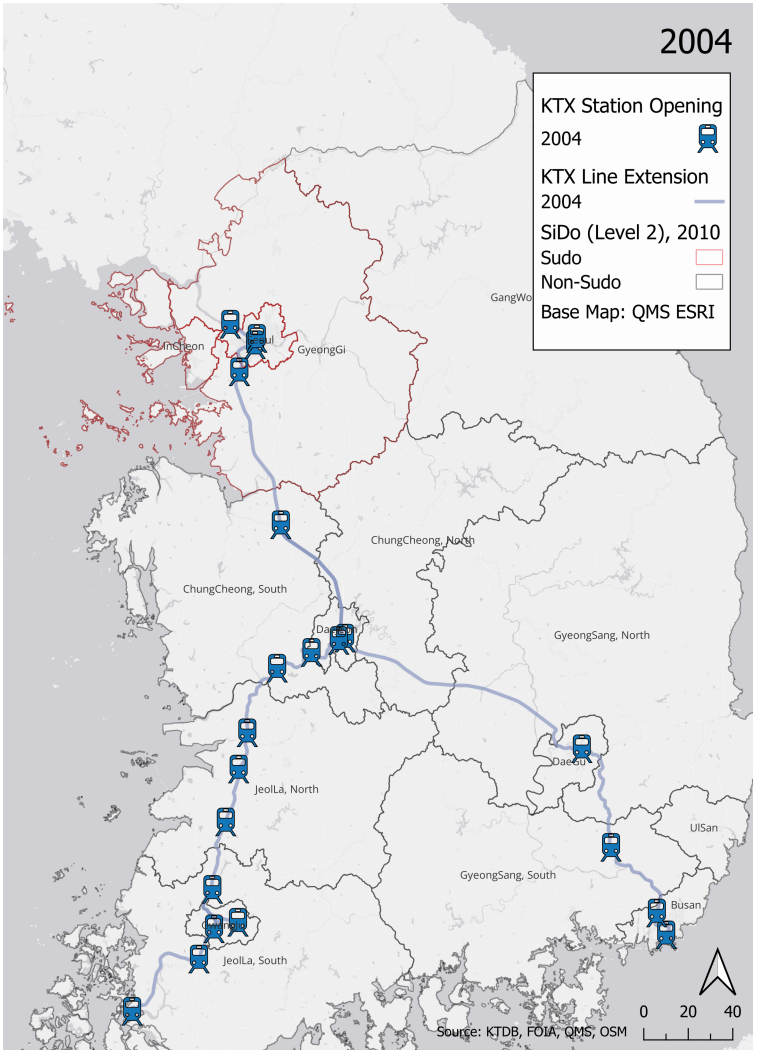
(d) Women (age 35-44)



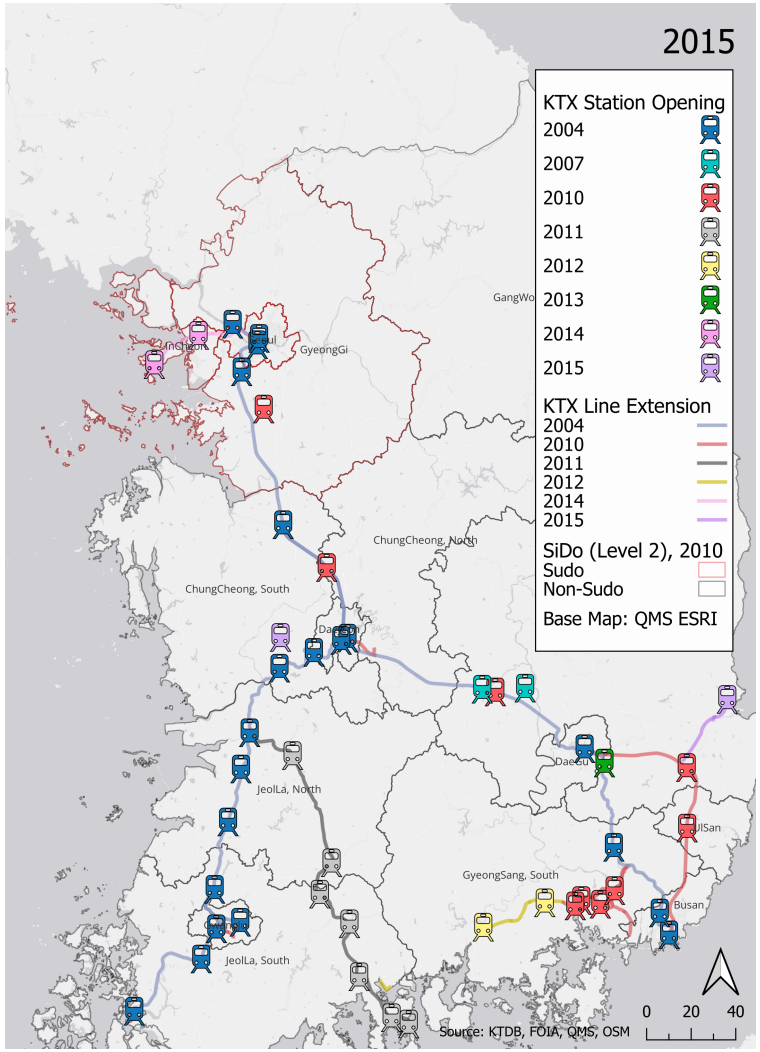
Notes: Each panel plots the occupational composition for a given gender and age group, based on the Resident Registration Population Survey (2010). Panels (a)–(b) show men aged 25–34 and 35–44; Panels (c)–(d) show women in the same age groups. Women aged 35–44 shift markedly toward sales and service occupations relative to their younger cohort, reflecting labor market re-entry after marriage and childbirth.

Figure A4: KTX Network Expansion

(a) 2004



(b) 2015



Notes: The map depicts the Korea Train eXpress (KTX) network in 2004 (Panel (a)) and in 2015 (Panel (b)). The red-bordered area represents the Seoul metropolitan areas. Source: Korea Transportation Database (KTDB).

B Data Appendix

B.1 Harmonization of Districts

During our analysis period from 2000 to 2015, there were many administrative changes. We harmonized the districts into their most consolidated form using the 2010 administrative boundaries as a guideline. After submerging districts into the most consolidated version and removing island districts not accessible via ground transport, our final analysis sample contains 228 districts. Different data sources use different administrative codes; we harmonized the MOIS code to the KOSIS version using the crosswalk file published by the Seoul Institute.

B.2 Census on Establishments

The Census on Establishments is an annual business survey covering all 3.3 million establishments in South Korea. It provides detailed business information including five-digit industry codes, geographic location, and employment counts by gender. Around five percent of the data contains missing worker counts, which we address by imputing from adjacent years' means.

B.3 Korean Labor and Income Panel Study (KLIPS)

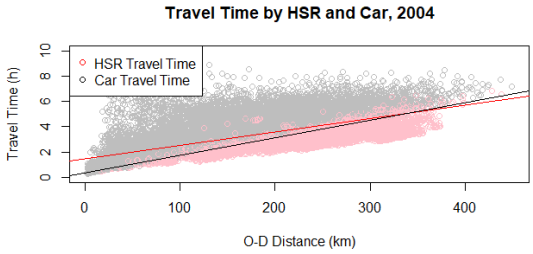
KLIPS is a comprehensive longitudinal survey of labor market activities. The initial 1998 sample extracted 5,000 households from urban districts; a second batch of 1,400 households from rural areas was added in 2009. We construct weighted average real wages at the district level using both cross-sectional and longitudinal weights. The small sample size (approximately 5,000 annually with complete wage and weight information) limits statistical power but provides the best available wage data for South Korea.

C Driving Time Calculation

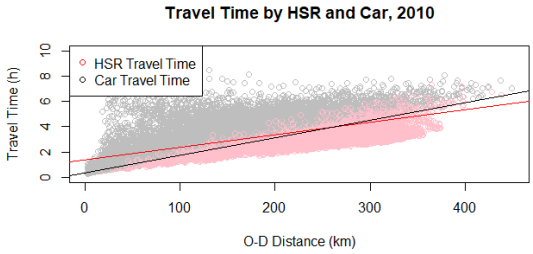
We calculate driving times from district centroids to the nearest KTX station using GIS network analysis from the Korea Transport DataBase (KTDB). OD travel time matrices are constructed for each year based on the KTX network available in that year. Figure A5 plots district-to-district travel times by KTX versus car for each year of the network expansion. Figures A6–A7 map the treated districts and the historical railroad buffer zones.

Figure A5: District to District Travel Time by KTX and Car

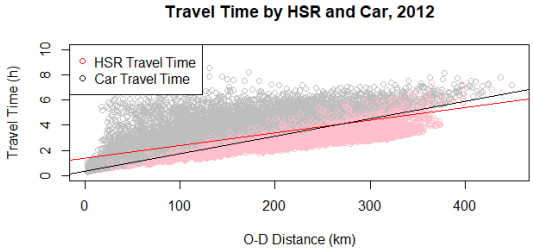
(a) 2004



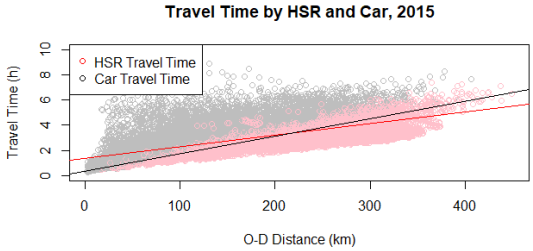
(b) 2010



(c) 2012

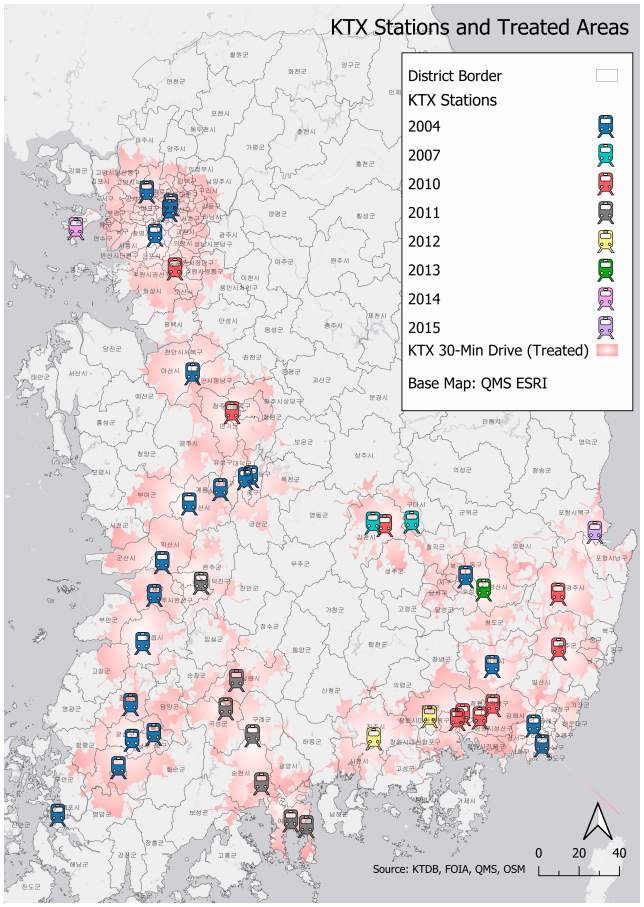


(d) 2015



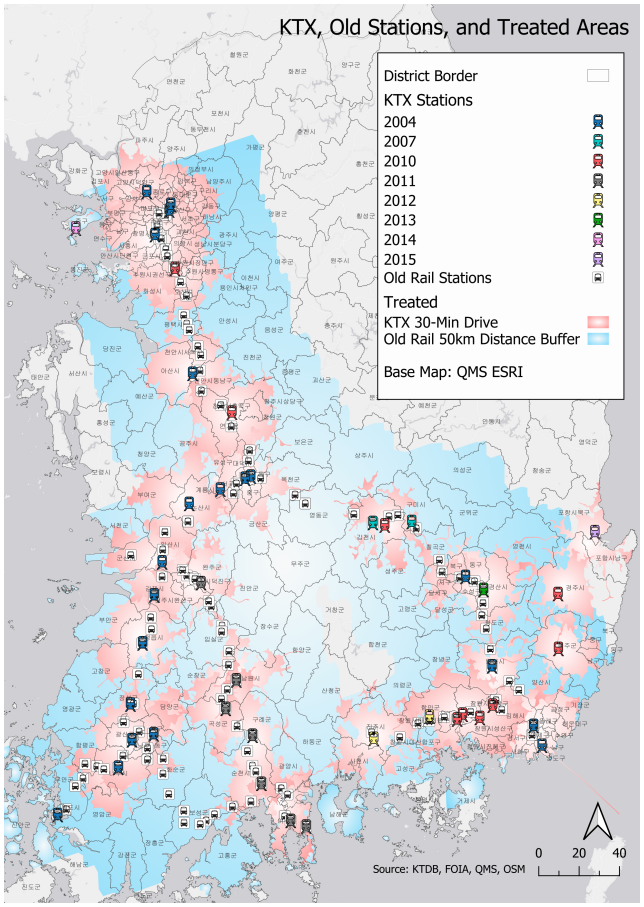
Notes: Each panel plots district-to-district travel times by KTX (red line) versus car (black line) for a given year of the KTX network. The horizontal axis ranks district pairs by car travel time; the vertical axis shows hours. Where the red line falls below the black line, KTX offers a faster connection. As the network expands from 2004 to 2015, KTX travel time advantages extend to a larger share of district pairs. Source: Korea Transportation DataBase (KTDB).

Figure A6: KTX and Treated Areas



Notes: The map shows the spatial distribution of KTX-treated districts as of 2015. A district is classified as treated if its geographic centroid falls within 30 minutes' driving time of the nearest open KTX station. Driving times are calculated using GIS shortest-path network analysis on the national road network. Shaded areas indicate treated districts; unshaded areas are never-treated or not-yet-treated controls. Source: Korea Transportation DataBase (KTDB).

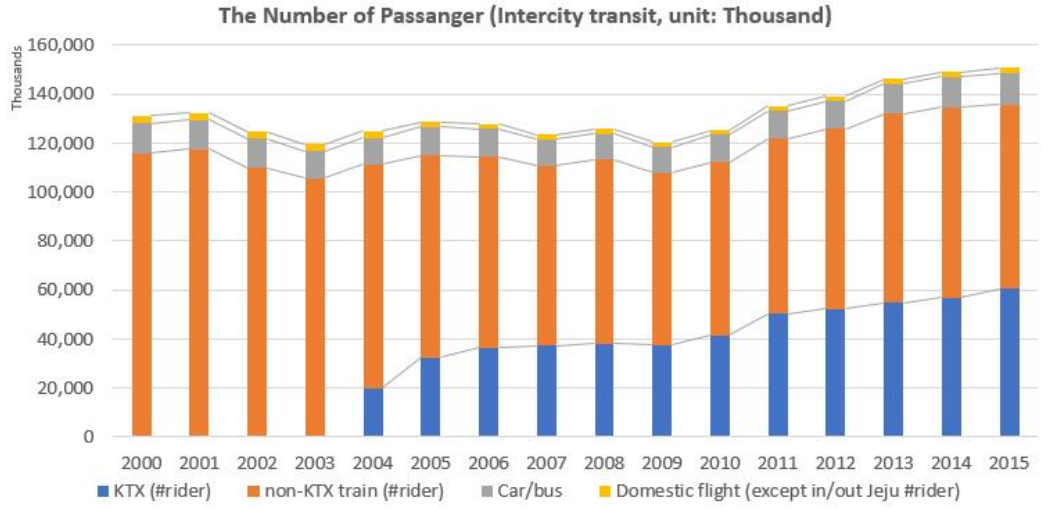
Figure A7: Overlap of KTX Treated Districts with Old Railroad Buffers



Notes: The map overlays KTX-treated districts (as defined in Figure A6) with buffer zones around railroad stations built during the Japanese colonial era (pre-1920). Brown-shaded areas indicate districts whose centroids lie within 50 km of an old railroad station. Our preferred specification restricts the analysis sample to districts within 80 km of these historical stations to ensure that treated and control districts share similar long-run development trajectories. Source: Korea Transportation DataBase (KTDB) and Korea Railroad Research Institute historical station records.

Figure A8: Annual Ridership across Different Transportation Modes

(a) Passenger Volume



(b) Total Passenger Travel Distance (KM)



Notes: Panel (a) shows annual passenger volume (millions) across intercity transportation modes in South Korea: KTX, conventional rail, express bus, and intercity bus. Panel (b) shows total passenger-kilometers traveled. KTX ridership grew rapidly after its launch in 2004, while conventional rail and intercity bus ridership declined. Source: Korea Transportation DataBase (KTDB).

D KTX Station Expansion History

Table A1: History of the KTX Station Expansion

Date	Line	From/To	Project Type	Station
2004 (+21)	Gyeongbu	Seoul/Daegu*	New line	Seoul, Hangshin, Gwangmeng, Cheonan-Asan, Daejeon, Dong-Daegu
		Daegu/Busan**	New line	Gupo, Milyang, Busan
	Honam**	Daejeon/Mokpo	New line	Yongsan, Seo-Daejeon, Gyeryong, Nonsan, Iksan, Gimje, Jeongeup, Jangsung, Gwangju, Gwangju-songjeong, Naju, Mokpo
2007 (+2)	Gyeongbu	Daegu/Busan	Add-station	Gumi, Gimcheon
2010 (+9, -2)	Gyeongbu*	Seoul/Daegu	Add-station	Suwon, Osong
		Daegu/Busan	Add/Close-station	Shin-Gyeongju, Ulsan, Gumi-Gimcheon, (Gumi), (Gimcheon)
			Improve speed	Busan
	Gyeongjeon**	Milyang/Masan	New line	Jinyoung, Changwon-Joongang, Changwon, Masan
2011 (+7)	Jeonla**	Ilsan/Yeosu	New line	Jeonju, Namwon, Guryegu, Sooncheon, Yeocheon, Yeosu-expo, Goksung
2012 (+2)	Gyeongjeon**	Masan/Jingu	Add-station	Haman, Jinju
2013 (+1)	Gyeongbu	Daegu/Busan	Add-station	Gyeongsan
2014 (+2)	Int'l Airport*	Incheon Airport/Seoul	New line	Gumam, Incheon International Airport
2015 (+1, -3)	Donghae*	Busan/Pohang	New line	Pohang
	Multi lines	-	Add/Close-station	Gongju, (Gimje), (Jangsung), (Haman)

Notes: *High Speed Railroad is defined to be maximum speed greater than $305km/h$. **Electrified conventional railway directly connected with HSR ($V_{max} < 180km/h$). Source: KORAIL.

E Industry Definitions and Regression Results

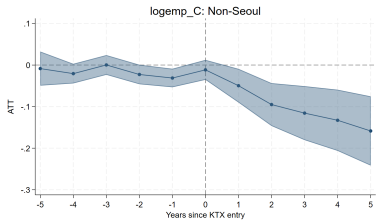
Table A2: Definition of Industry

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	F-to-M ratio		Seoul		Non Seoul		
Industry Name	Seoul	Non Seoul	Emp.	Share	Emp.	Share	Included KSIC9 (3 digit)
Transportation and Logistics	0.120	0.111	465,346	0.080	423,316	0.070	all H
Construction, Utility	0.149	0.134	308,205	0.053	335,643	0.055	all D and F
Government	0.297	0.251	130,292	0.022	190,816	0.031	O841-O845
Wholesales	0.374	0.427	544,876	0.094	544,876	0.090	G451-G453, G461-G468
Other Services	0.376	0.408	186,875	0.032	223,893	0.037	S941, S942, S949, S951, S952, S953
Real Estate and Rental Service	0.416	0.413	239,009	0.041	139,470	0.023	all L but L694
Manufacturing	0.624	0.709	1,386,472	0.239	1,386,472	0.228	all C but C191, E383
Other Retail	0.974	0.973	357,417	0.062	421,095	0.069	G473, G475-G479
Lodging	1.117	1.564	55,661	0.010	80,423	0.013	I551, I559
Food, Bev, Tobacco Retail	1.173	1.386	81,990	0.014	118,622	0.020	G472
Other Private Education	1.226	1.177	333,068	0.057	326,881	0.054	P854-P856, R902, R911, R912
General Retail	1.236	1.453	167,172	0.029	185,463	0.031	G471
Banking, Finance, Insurance	1.500	1.229	286,050	0.049	236,073	0.039	all K and L694
Restaurant and Bar	2.157	2.587	733,845	0.126	852,016	0.140	I561, I562
Medical	2.192	2.059	212,860	0.037	221,073	0.036	Q861-Q863, Q869
Textiles, Clothing Retail	2.303	2.660	112,989	0.019	126,384	0.021	G474
Hair Salon and Spa	2.764	2.407	100,799	0.017	113,831	0.019	S961
Pre and Elementary School	3.825	2.356	107,898	0.019	143,831	0.024	P851

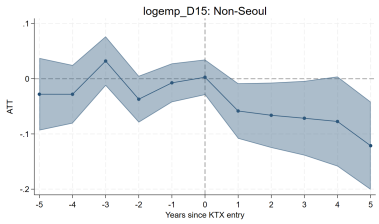
The table lists the industry names used in our analysis and the 3-digit level industry codes included in each category. The categorization is based on the industry codebook of the Census on Establishments. The list is sorted based on column (2). Table 2 uses the following sector mappings from this table: Construction & Utilities = Construction, Utility; Banking & Finance = Banking, Finance, Insurance; Retail = General Retail; Education = Pre and Elementary School + Other Private Education; Healthcare = Medical. Arts & Recreation (KSIC R) and Business & Professional (KSIC M744), included in Table 2, are not separately listed here.

Figure A9: Sector-Specific Event Study Plots (Non-Seoul)

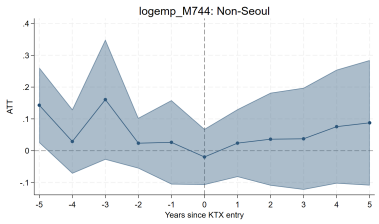
(a) Manufacturing



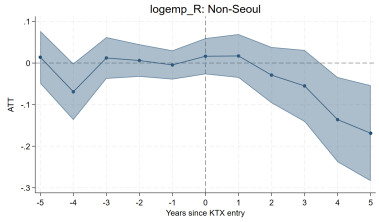
(b) Construction & Utilities



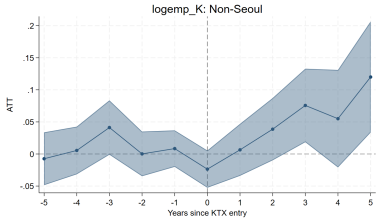
(c) Business & Professional



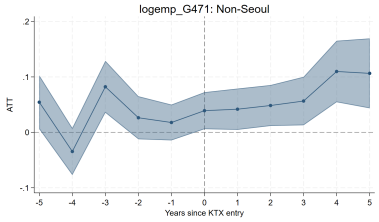
(d) Arts & Recreation



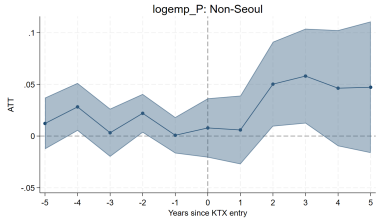
(e) Banking & Finance



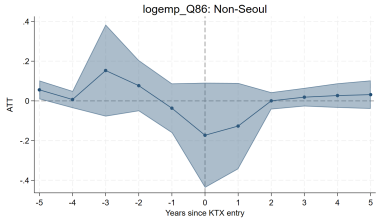
(f) Retail



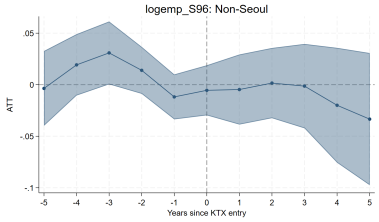
(g) Education



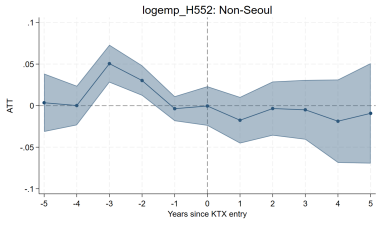
(h) Healthcare



(i) Hair Salon & Spa

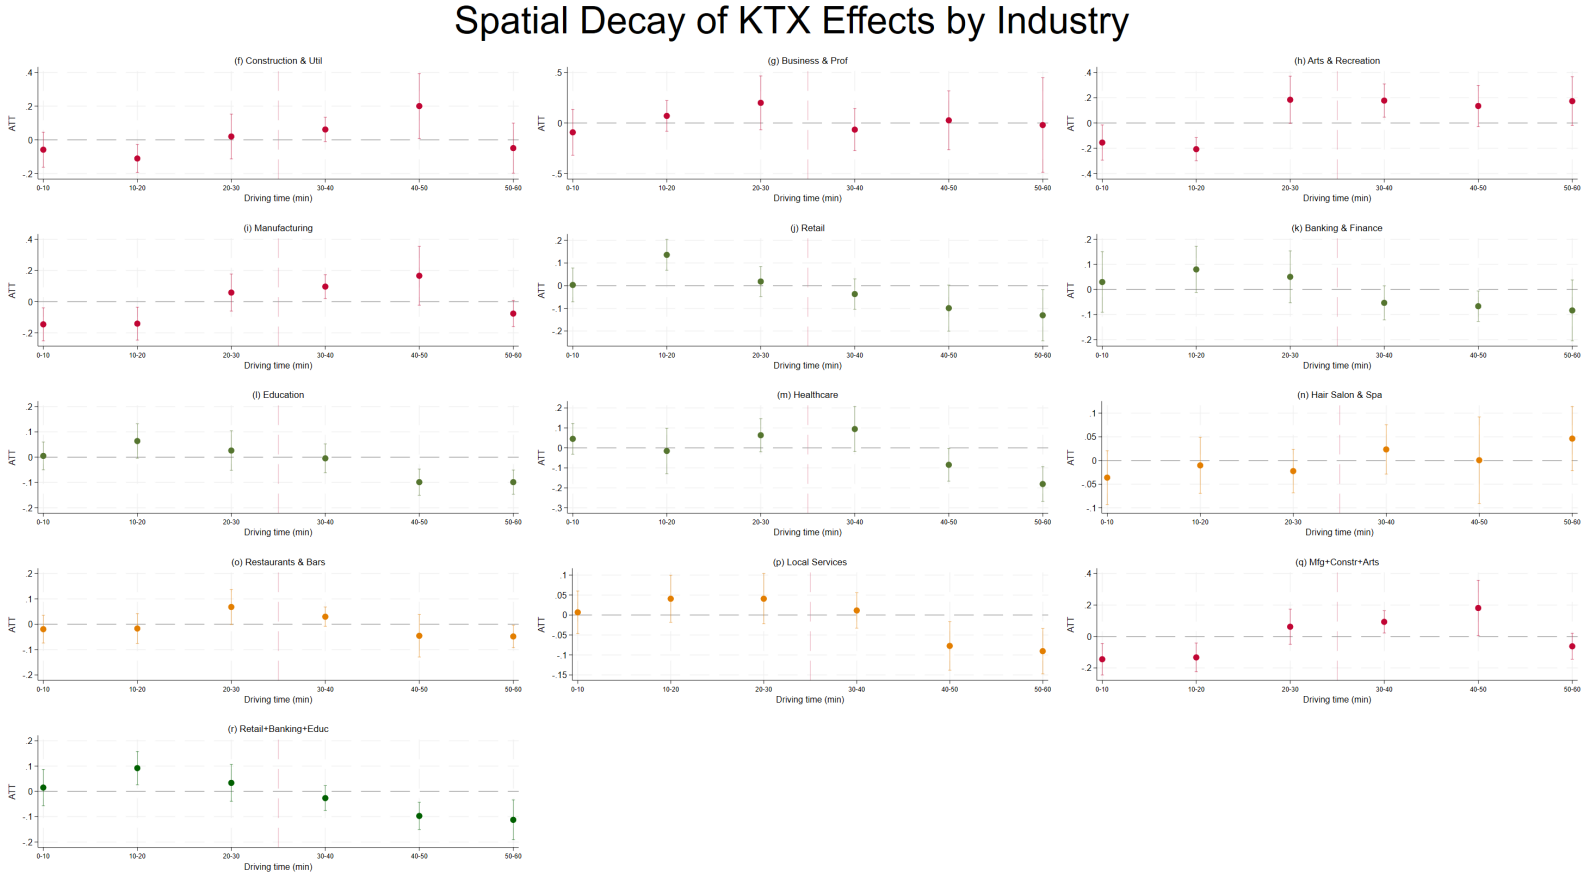


(j) Restaurants & Bars



Notes: Each panel shows the staggered DiD event-study coefficients for a sector's log employment. The horizontal axis denotes years relative to KTX entry. Panels are ordered by baseline (2003) female employment share, ascending (male-intensive at top, female-intensive at bottom). Dashed horizontal line marks zero. 95% confidence intervals shown; standard errors clustered at the district level. Sample: Non-Seoul districts within 80 km of historical railroad stations. See Table 2 for aggregated ATT estimates and Table A2 for sector definitions.

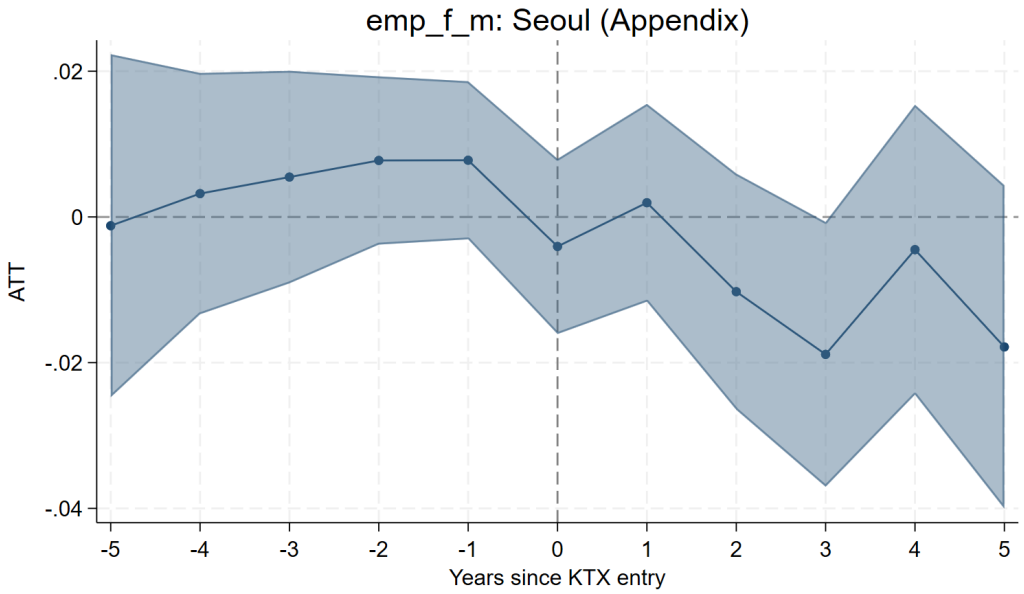
Figure A10: Spatial Decay of KTX Effects by Industry



Notes: Each panel plots staggered DiD ATT estimates by driving-time ring (10-minute bins) for an individual industry. For each ring, treated districts are those within the ring; controls are all other districts within the 80 km colonial-railway buffer (not-yet-treated as comparison). Bars represent 95% confidence intervals. Standard errors clustered at the district level. Aggregate and gender composition outcomes are shown in Figure 3.

F Seoul Placebo Results

Figure A11: Seoul Placebo: Event Study Plots



Notes: Staggered DiD event study plots for Seoul metropolitan districts. The statistically insignificant emp_f_m result (ATT= -0.009) confirms the gender employment gap effect is specific to non-Seoul areas.