Revisiting the Fertility Puzzle:

A Dynamic Spatial Equilibrium Approach*

Costas Arkolakis† Jun Hee Kwak‡ Jaemin Woo§ Hyunjoo Yang¶

Yale University & NBER Sogang University Brown University Sogang University

March 2026

Preliminary and incomplete. Please do not circulate.

Abstract

We study fertility in a dynamic spatial equilibrium model with endogenous location and residential choices, alongside investments in children's education. We estimate the model using near-universal administrative microdata from South Korea, drawing on census records, credit bureau data linked to mobile phone mobility records, and government administrative sources. Our estimates indicate that high commuting and housing costs, combined with intensive education investments in large cities, where the amenity of living is large, substantially reduce fertility. We conduct a decomposition exercise to explain the sharp decline in South Korean fertility over the past decade, finding that deteriorating child-rearing amenities and rising education costs play a significant role. Finally, we conduct a series of counterfactual exercises considering place-based policies to reduce commuting costs, increase housing supply, or lower education costs. We find that policies focusing exclusively on large cities, such as Seoul, may have the opposite effect on fertility than intended due to spatial sorting. However, a combination of nationwide policies targeting the fertility margin directly can boost fertility significantly.

*We thank Siyoung Lee for excellent research assistance. The authors have no relevant financial relationships or conflicts of interest to disclose.

†costas.arkolakis@yale.edu

‡jhkwak@sogang.ac.kr

§jaemin_woo@brown.edu

¶hyang@sogang.ac.kr

1 Introduction

Fertility has fallen sharply across advanced economies, reshaping labor supply, the age structure of the population, and the fiscal sustainability of social insurance systems. South Korea represents the extreme case. According to Statistics Korea, the total fertility rate reached 0.72 in 2023 and rose only slightly to about 0.75 in 2024, remaining the lowest in the world (Statistics Korea, 2024). What explains these persistently low birth rates? A large body of work emphasizes changes in gender roles, social norms, labor market frictions, and the difficulty of combining careers and family life in modern economies (e.g., Doepke et al., 2023; Goldin, 2021; Kearney and Levine, 2025). These explanations are important, yet they are often silent about a striking empirical regularity. Fertility varies enormously across space within a country, and the lowest fertility is frequently concentrated in dense metropolitan areas.

This paper argues that fertility should be analyzed as the equilibrium outcome of a dynamic spatial household problem. Households choose where to live, where to work, how much housing to consume, and how much to invest in their children's human capital. These choices interact in general equilibrium. Spatial sorting affects housing demand, commuting patterns, and local prices. It also changes the intensity of competition for educational status and the private cost of raising children. In this environment, fertility is not merely a preference parameter or a reduced-form outcome. Rather, it is an endogenous choice that responds to local wages, housing costs, commuting costs, and the education environment.

We make three contributions. First, we develop a dynamic spatial equilibrium model that integrates fertility, location choice, commuting, and education investment in a unified framework with market clearing. Unlike models that treat fertility or location in isolation, our framework captures how sorting, housing prices, and education competition interact in general equilibrium. Second, we assemble and link administrative microdata covering nearly the entire Korean working-age population, providing a detailed empirical characterization of the spatial fertility gradient and its correlates at national scale. Third, we leverage the estimated model to decompose the drivers of the fertility decline and to evaluate the efficacy of various place-based policies.

We bring this framework to data using a combination of administrative and survey sources from South Korea. Migration flows and birth outcomes come from the Population and Housing Census.

The Korea Credit Bureau (KCB) database covers the near-universe of the working-age population, approximately 40 million individuals, and can be linked to mobile phone mobility data from SK Telecom (SKT) within a secure processing environment. This linked data environment, enabled by landmark 2021 legislation permitting secure individual-level linkage across administrative systems, has only recently become available for economics research and has begun to be used in other work (Arkolakis et al., 2025). The SKT mobility data records the dong-level commute times of individuals utilized in our descriptive analysis and provides the foundation for modeling commuting decisions. Additionally, the Census and KCB datasets provide the commuting flow shares that enter the structural estimation of commuting costs. We complement these sources with government data on housing prices, wages, education expenditures, and local amenities to discipline the model and identify the remaining structural parameters.

The model's parameters are identified through a sequential estimation strategy that exploits the discrete choice structure. Moving costs are recovered from a gravity equation estimated on bilateral migration flows, with distance and province boundaries as shifters. Commuting costs are identified from the ratio of commuting flows across workplace destinations for workers in the same residential location, using a Bartik-style instrument based on industry composition to address the endogeneity of wages. Education costs are estimated from the Private Education Survey by regressing expenditure shares on location, parental education, and number of children. The remaining preference parameters, including local amenities and the fixed cost of children, are recovered as structural residuals from birth rate regressions that condition on the estimated cost parameters. Section 5 provides details.

Empirically, we document several central stylized facts. Across neighborhoods, fertility is strongly negatively correlated with population density, housing prices, and commuting costs, while it is positively associated with amenity measures such as childcare availability. High-skilled workers exhibit substantial net inflows into the Seoul metropolitan area, whereas low-skilled workers show much weaker spatial sorting. Education costs are highest in Seoul and rise steeply with the number of children, consistent with increasing marginal costs of child quality in high-competition environments. Taken together, these facts suggest that high urban amenities are insufficient to offset the compounded costs of raising children in dense urban centers.

Our quantitative results show that the interaction of spatial sorting, education competition, and

general equilibrium housing and commuting responses can generate large fertility differences across space. The model helps explain why fertility remains low even in rich and amenity-abundant cities. Urban areas offer high wages and amenities, which raise the value of children, but these benefits are outweighed by the equilibrium increase in the cost of children through housing, commuting, and education expenditures. The resulting quantity-quality trade-off implies that standard policies targeted at only one margin may deliver limited welfare gains when deep spatial frictions and competitive externalities remain in place.

Our structural decomposition reveals that the decline in South Korea's completed fertility from 2.0 to 1.5 between 2005 and 2020 was primarily driven by within-region behavioral responses to shifting preferences and rising costs. While deteriorating child-rearing amenities and rising education costs are the dominant factor (explaining around 23 percent of the decline), the preference shock also play a substantial role (accounting for a further 16 percent). Interestingly, we find that spatial sorting—the compositional effect—partially offsets these declines for certain factors like housing, though this reallocation is insufficient to counter the broader downward trend.

We also use the estimated model to evaluate fertility policies, and the counterfactual experiments yield three main findings. First, place-based policies that target the Seoul metropolitan area alone can backfire. Reducing commuting costs, expanding housing supply, improving amenities, or raising productivity in Seoul all produce positive within-region fertility effects, but they also attract households into a region where baseline fertility is low. The resulting compositional shift offsets the behavioral gain. A commuting cost reduction equivalent to 20 percent of consumption in Seoul lowers national fertility by 2.7 percent, and doubling Seoul's housing stock reduces the fertility by 1.8 percent, as it attracts people to Seoul with a lowest fertility rate.[1]

Second, education cost reduction is the most powerful single lever. A nationwide education cost reduction equivalent to 20 percent of consumption in welfare terms raises the national fertility rate by 16.2 percent, several times larger than any other single-margin intervention, because it directly weakens the quantity-quality wedge that spatial sorting amplifies.

Third, coordinated policies outperform isolated interventions. Subsidies to parents financed by a 20 percent tax on nonparents raise fertility by 8.8 percent on their own, a modest effect in line with reduced-form estimates of cash transfer programs (Milligan, 2005; González, 2013). But

[1]All percentage changes are reported as log differences, $100 \times \ln(\text{TFR}_{cf}/\text{TFR}_{base})$.

combining these subsidies with education reform raises fertility by 21.7 percent, roughly the sum of the two individual effects, because the policies target largely distinct margins and their effects are approximately additive. These results suggest one reason why Korea's sustained pronatalist spending has not reversed the fertility decline, and why interventions targeting the structural source of low fertility may achieve substantially more than transfers alone.

1.1 Related Literature

This paper draws on and contributes to work in fertility economics, housing and family formation, quantitative spatial equilibrium, and education competition.

The modern economics of fertility emphasizes gender roles, career-family compatibility, and social norms as drivers of declining birth rates (Becker, 1960; Becker and Tomes, 1976; Barro and Becker, 1989; Doepke et al., 2023; Doepke and Kindermann, 2019; Goldin, 2021; Kearney and Levine, 2025). These frameworks do not model location choice. They explain why fertility has fallen on average but not why it varies so sharply across locations within the same country, or why the highest-wage cities often have the lowest fertility. We show that the same household may choose different fertility in Seoul versus a rural province because local prices, commuting costs, and education competition differ. This spatial margin is quantitatively important and has distinct policy implications.

A growing literature links housing costs to fertility. Reduced-form studies find that house-price increases raise fertility for homeowners through wealth effects, while effects for non-owners are small or negative (Dettling and Kearney, 2014; Lovenheim and Mumford, 2013; Ang et al., 2024; Fazio et al., 2024). Couillard (2025) estimates that undoing the post-1990 rise in U.S. housing costs would increase fertility by 11%, accounting for roughly half the TFR decline. Many of these reduced-form studies treat local housing costs as exogenous shifters. We embed housing in a spatial equilibrium where prices respond endogenously to sorting. When high-skilled workers move to Seoul for higher wages, they bid up housing prices as a consequence of the same forces that raise their earnings. The observed negative correlation between prices and fertility thus conflates a compositional effect, where low-fertility households sort into expensive locations, with any behavioral response to prices. Our framework separates these channels and shows that cross-sectional patterns are consistent with

a large compositional component.[2]

We build on quantitative spatial equilibrium models with discrete location choice and market clearing (Diamond, 2016; Redding and Rossi-Hansberg, 2017; Allen and Arkolakis, 2014, 2025) and dynamic spatial frameworks with forward-looking migration (Kleinman et al., 2023; Eckert and Kleineberg, 2024). Many quantitative spatial equilibrium models focus on labor-market outcomes and abstract from fertility. Some recent work (e.g., Eckert and Kleineberg, 2024) incorporates forward-looking education and residential choices, but fertility and parental child-investment decisions are typically absent. We add a quantity-quality tradeoff where the local education environment shapes the marginal cost of children. We also add an intergenerational channel where parents' location choices affect children's future opportunities. These extensions introduce new channels through which sorting affects welfare. A policy that attracts high-skilled workers to a city raises local wages but also intensifies education competition, potentially lowering fertility and well-being for incumbent residents.

Education competition is central to our mechanism and particularly salient in East Asia. Kim et al. (2024) show that status externalities in education substantially lower Korean fertility, with eliminating the externality raising fertility by 28% in their model. Keng et al. (2011) find that a Taiwanese reform reducing education subsidies lowered fertility by 17% among public-sector households. These papers study education externalities in representative-agent or symmetric-location models where competition intensity is uniform across space. We endogenize competition through spatial sorting. When high-skilled parents cluster in Seoul, they raise the local benchmark for educational achievement, and all parents in Seoul respond by investing more heavily. This spatial heterogeneity in competition is what allows place-based policies to affect fertility. A childcare subsidy in Seoul attracts families, intensifies the local rat race, and may partially offset its own effect on birth rates.

Finally, our counterfactual exercises contribute to the policy evaluation literature. Reduced-form studies of cash transfers, childcare subsidies, and parental leave typically find modest fertility effects, often in the single digits and sometimes affecting only timing rather than completed fertility.[3] In Korea, Kim (2024) finds that municipal baby bonuses raised the TFR by about 4.7%.

[2]Earlier frameworks modeling housing and fertility jointly include Coeurdacier et al. (2023) and Simon and Tamura (2009).

[3]On cash transfers, see Milligan (2005), González (2013), Cohen et al. (2013), Gauthier and Hatzius (1997),

These estimates, however, do not account for the migration, housing price, and education competition responses that our model shows adjust endogenously to policy. Structural models that incorporate general equilibrium feedbacks can find larger effects when multiple margins adjust simultaneously (Fehr and Ujhelyiova, 2013; Guner et al., 2024), though results depend on which margins are modeled (Table A.5). Our contribution is to evaluate place-based policies through a spatial lens, showing that interventions in one location spill over to others through migration and price adjustments.

1.2 Roadmap

The remainder of the paper proceeds as follows. Section 2 describes our main data. See Appendix A for the institutional setting. Section 3 documents key spatial facts about fertility, education spending, housing prices, and commuting costs. Section 4 presents the dynamic spatial equilibrium model and its mechanics. Section 5 outlines the estimation strategy and identification. Section 6decomposes factors explaining the decline in fertility rates. Section 7 conducts counterfactual policy experiments, decomposing each policy's national fertility effect into within-region behavioral changes and cross-region compositional shifts. Section 8 concludes.

2 Data

The empirical analysis requires joint observation of residential location, workplace, fertility outcomes, and education spending at national scale. Korea's administrative data infrastructure, enabled by landmark 2021 legislation permitting secure individual-level linkage across government systems, makes this possible.

The KCB-SKT linked mobility data provide two inputs to the analysis. Commuting flow shares between residential and workplace locations enter the structural estimation of commuting costs described in Section 5. Dong-level commute times serve as covariates in the descriptive regressions in Table 1. Migration flows and birth shares by education level and province are constructed from the

Slonimczyk and Yurko (2014), and Bokun (2024). On childcare, see Bauernschuster et al. (2016), Kravdal (1996), Rindfuss et al. (2010), and Havnes and Mogstad (2011). On parental leave, see Lalive and Zweimüller (2009), Raute (2019), Dahl et al. (2016), Hoem (1993), and Laroque and Salanié (2014). Gauthier (2007) and Doepke et al. (2023) provide comprehensive reviews. Appendix G compiles these estimates systematically.

Population and Housing Census. A worker is classified as high-skilled if they have graduated from a four-year university. We use "college-educated" and "high-skilled" interchangeably throughout.

We complement these records with several external sources. Fertility outcomes come from the Population and Housing Census and vital statistics published by Statistics Korea. Housing prices come from apartment transaction records published by the Ministry of Land, Infrastructure, and Transport. Education costs come from the Private Education Expenditures Survey, jointly administered by Statistics Korea and the Ministry of Education.[4]

The analysis uses three geographic levels. The coarsest is the *si-do*, Korea's 17 provinces and metropolitan cities.[5] The intermediate level is the *si-gun-gu*, roughly equivalent to a U.S. county, of which there are approximately 250. The finest level is the *dong*, an administrative neighborhood of which there are approximately 3,500 nationwide. The descriptive regressions in Table 1 exploit dong-level variation within si-gun-gu. The structural model operates at the si-do level, primarily for computational tractability. The state space of residence-workplace pairs grows quadratically in the number of locations, and solving the model at the si-gun-gu level would require optimizing over roughly 62,500 pairs rather than 289.[6] The si-do level nonetheless captures the key spatial margins in the data, including the large wage and housing price gaps between Seoul and other provinces, the cross-province commuting patterns that allow households to separate residence from workplace, and the provincial variation in education competition that the model formalizes.

3 Stylized Facts and Empirical Motivation

This section uses our linked administrative data to establish five empirical patterns that motivate the spatial equilibrium framework developed in Section 4. The patterns share a common structure. Fertility varies dramatically across Korean provinces, and this variation is strongly correlated with housing costs, commuting burdens, and the intensity of education competition. Yet high-wage

[4]Wages and labor supply at the province level come from the Local Area Labor Force Survey, accessed through Statistics Korea's Microdata Integrated Service (MDIS). Housing expenditure shares come from the Housing Survey. Population counts by age and province come from the Ministry of the Interior and Safety. Appendix B describes the construction of each variable in detail.

[5]The 17 si-do units include Seoul, six metropolitan cities (Busan, Daegu, Incheon, Gwangju, Daejeon, Ulsan), Sejong Special Autonomous City, and nine provinces.

[6]Future drafts will include robustness analyses at more disaggregated geographic levels, including si-gun-gu-level specifications, to verify that the main quantitative results are not driven by within-province heterogeneity that the si-do-level model does not capture.

locations are precisely where fertility is lowest. This juxtaposition of high wages and low fertility is difficult to explain with standard models in which income raises the demand for children. What distinguishes our setting is that high wages attract workers, and the resulting concentration of high-skilled households raises housing prices and intensifies education competition, creating a spatial trade-off that erodes the income gains.

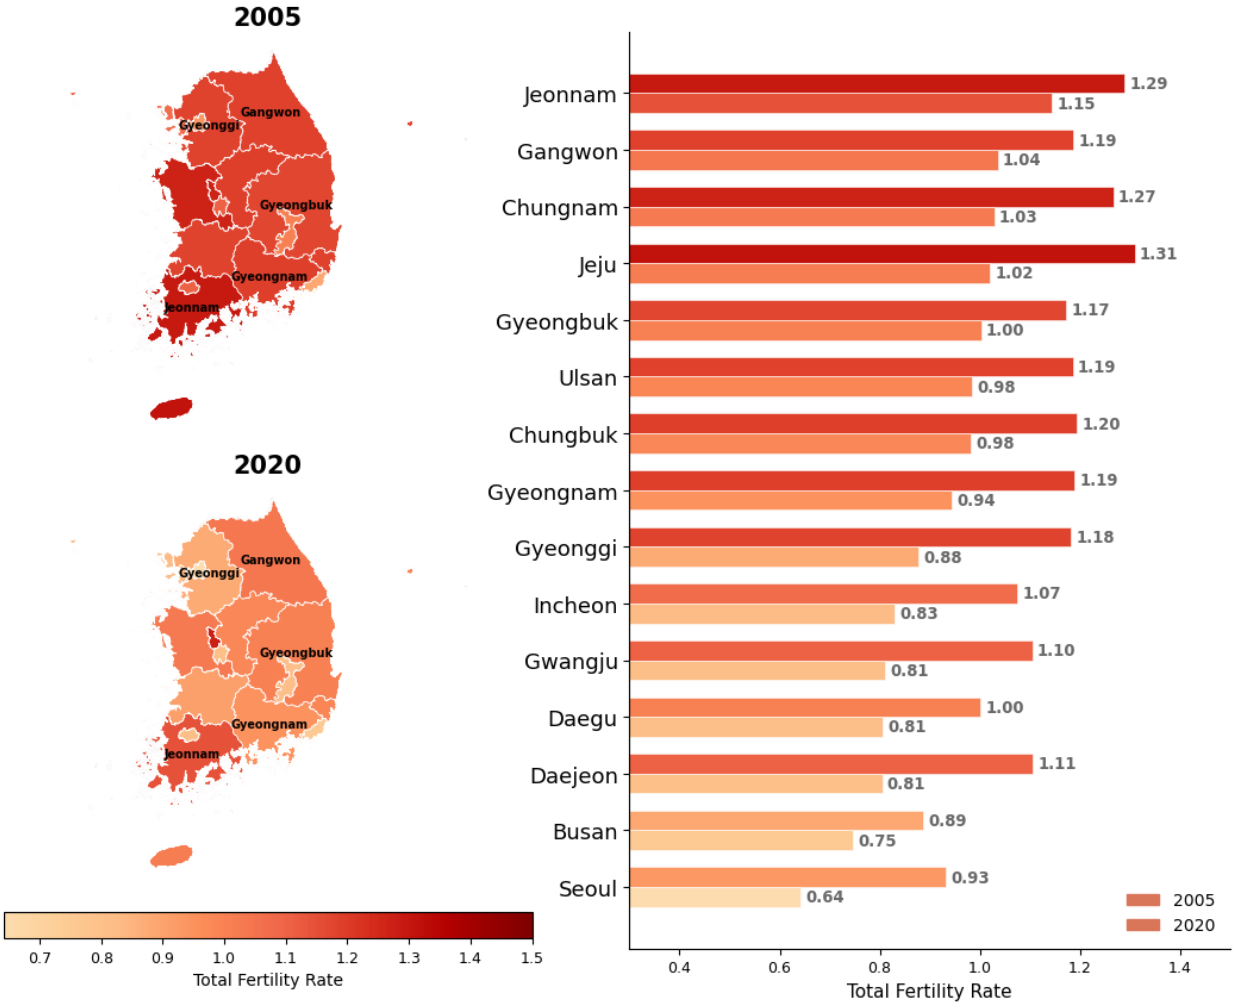
Three fertility measures appear throughout the paper, each suited to a different level of analysis. The *total fertility rate* (TFR), defined as the average number of children a woman would bear over her lifetime given current age-specific birth rates, summarizes aggregate fertility at the province level and is the metric used in the introduction. The *birth rate*, defined as births per 1,000 women aged 20–44, provides a cross-sectional measure at the dong or si-gun-gu level and is the dependent variable in Table 1. In the structural model, fertility is measured as the total number of children born divided by the number of young adults (aged 40–50), a completed cohort fertility measure that differs in level from the period TFR. Section 7 discusses the distinction and reports percentage changes as the primary outcome. The spatial patterns documented below are robust across all three measures.

The model's baseline fertility rate of 1.5 exceeds the period TFR of 0.72 for two reasons. First, the period TFR is a synthetic cohort measure that sums current age-specific birth rates. When women delay childbearing, the period TFR falls below the number of children women eventually have over their lifetimes. Cohort completed fertility for Korean women born in the 1970s is approximately 1.5 to 1.7 (Hwang, 2023), much closer to the model's baseline. Second, the model operates with a coarse OLG structure in which a single "young adult" period spans ages 25 to 55, so its fertility measure captures lifetime births within that window rather than a single calendar year's flow. Because the counterfactual experiments hold this structure fixed and compare equilibria under identical definitions, percentage changes in the model's fertility measure remain informative about the direction and relative magnitude of policy effects even though the level differs from the period TFR.

3.1 The Geographic Fertility Gradient

South Korea's aggregate fertility rate of 0.72 conceals enormous geographic dispersion. Figure 1 displays total fertility rates by province in 2023. Seoul records the lowest fertility in the country

Figure 1: Total fertility rate by province, 2023



Notes: Total fertility rate (TFR) by si-do for 2023, calculated from vital statistics records published by Statistics Korea. National TFR in 2023 was 0.72. See Appendix B for variable construction.

at 0.55, while Jeonnam and Sejong record the highest at 0.97, nearly twice Seoul's rate. The Seoul Metropolitan Area, home to roughly half the national population, sits well below the national average.

The coexistence of an urban fertility deficit with higher fertility elsewhere defines the core puzzle. Seoul offers the highest wages in the country, yet fertility there is lowest. A natural hypothesis is that households in Seoul simply have fewer resources available for children after paying high housing costs. But this explanation runs into difficulty when we examine within-location patterns. Table 1 shows that income and assets are positively associated with birth rates once we condition on location. With si-gu fixed effects, a one-decile increase in income predicts a 0.20 decile increase

in birth rates, and a one-decile increase in assets predicts a 0.43 decile increase. Resources help households afford children within a given location. The puzzle is that the locations where resources are most abundant are also the locations where fertility is lowest. The resolution we propose is that spatial variation in fertility reflects sorting and endogenous local prices operating in general equilibrium.

Table 1: Correlates of Birth Rates

	(1) Birth Rate	(2) Birth Rate	(3) Birth Rate
Income	0.325***	0.295***	0.198***
	(0.0400)	(0.0401)	(0.0462)
Total assets	0.261***	0.280***	0.433***
	(0.0724)	(0.0714)	(0.0758)
House price	-0.225***	-0.0622	0.112
	(0.0514)	(0.0657)	(0.0738)
Population density	-0.761***	-0.667***	-0.498***
	(0.0676)	(0.0791)	(0.0813)
Daycare density	0.464***	0.390***	0.415***
	(0.0601)	(0.0691)	(0.0738)
Education achievement	-0.145***	-0.173***	-0.175***
	(0.0246)	(0.0276)	(0.0277)
Commute time (male)	-0.103***	-0.101***	-0.0626**
	(0.0320)	(0.0310)	(0.0288)
Commute time (female)	-0.0101	-0.00744	-0.0128
	(0.0283)	(0.0283)	(0.0277)
Observations	1547	1547	1522
R-squared	0.287	0.322	0.523
Fixed Effects	None	SiDo Fe	SiGu FE

Notes: All variables are measured in deciles across dong (administrative neighborhoods), pooled and averaged over July 2019–May 2024 (Commute time is over Aug 2022–Dec 2023). Education achievement is measured using data from 2024 only. Birth rate is defined as births per 1,000 women aged 20–44. Deciles are constructed by ranking dong and grouping them into ten equally sized bins. Regressions are unweighted. Standard errors in parentheses are heteroskedasticity-robust. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3.2 Housing and Commuting Costs

Sorting translates into costs through two channels, housing and commuting. Table 1 confirms that population density remains strongly negative even with si-gu fixed effects, with a coefficient of -0.50. Housing prices, by contrast, lose statistical significance once fixed effects are included, falling from -0.23 without controls to $+0.11$ with si-gu fixed effects. This pattern reflects the

distinction between cross-sectional sorting and within-location behavioral responses. In the cross-section, high prices and low fertility both result from high-skilled workers sorting into desirable locations. The negative correlation reflects this composition effect rather than reflecting a direct causal impact of prices on fertility. Within a location, price variation is smaller and less clearly related to fertility decisions. The model incorporates both channels. Prices affect fertility through the budget constraint, but in cross-sectional data the dominant source of spatial variation appears to be sorting rather than the direct price effect. The attenuation in reduced-form coefficients does not imply that housing costs are irrelevant for fertility. Rather, it reflects sorting confounding cross-sectional comparisons. Our structural estimates in Section 5 isolate the behavioral response.

The second channel is commuting. Male commute time is negatively associated with fertility in the cross-section, and the relationship survives geographic controls with a coefficient of -0.06 after including si-gu fixed effects. Commuting time matters for fertility because it represents a direct drain on household time budgets. Long commutes reduce the time available for childcare, increase the need for outside childcare arrangements, and raise the effective cost of having children. The model formalizes this channel by treating commuting and migration as distinct margins, allowing households to access high-wage employment through longer commutes without necessarily relocating.

3.3 Education Competition

A distinguishing feature of our framework is the role of education competition, which operates as a separate wedge on fertility beyond housing and commuting costs. Table 1 confirms that this relationship survives both si-do and si-gu fixed effects with a coefficient of -0.17. We use academic achievement as a proxy for the intensity of education competition, on the grounds that districts where parents invest more heavily in private tutoring tend to produce higher test scores. Achievement thus reflects the competitive investment environment rather than measuring competition directly. Unlike housing prices, the education correlation does not disappear when controlling for geography. Districts with intense education competition have lower fertility even relative to nearby districts with similar density and price levels.

This finding motivates our treatment of education as an explicit choice margin in the model. When high-skilled households concentrate in a location, they raise the local competitive benchmark

for educational outcomes. Parents respond by investing more heavily in each child's human capital, but this investment is costly. The marginal cost of child quality rises with the intensity of local competition, sharpening the quantity-quality tradeoff and lowering fertility. This mechanism generates spatial variation in fertility that cannot be attributed to housing or commuting costs alone.

3.4 Child-Related Amenities

Childcare availability is positively associated with birth rates and is robust across specifications. Table 1 shows that daycare density carries a coefficient of 0.42 even with si-gu fixed effects. This finding suggests that childcare infrastructure lowers the cost of children by reducing the time burden on parents. Yet the positive effect of childcare does not overturn the urban fertility penalty. Locations with abundant childcare also tend to bundle high housing costs, long commutes, and intense education competition. The net effect in high-amenity urban locations is still lower fertility relative to less amenity-rich but cheaper alternatives.

3.5 Interpretation

A unifying interpretation of the patterns documented above is that they reflect sorting. High-skilled workers move to Seoul in pursuit of higher wages. This concentration raises housing demand and commuting congestion, pushing up the cost of living in the capital region. It also brings together parents who invest heavily in education, intensifying local competition and raising the marginal cost of child quality. Childcare amenities partially offset these costs but do not eliminate them. The result is that urban areas offer high wages and high amenities but also high costs, and the net effect on fertility is negative.

These correlations are descriptive and do not establish causal relationships. High-cost locations may attract households with lower unobserved fertility preferences, and the negative correlation between achievement and fertility could reflect reverse causality if fewer children allow more investment per child. The structural model developed in Section 4 addresses these concerns by explicitly modeling the joint choice of location and fertility, separating the causal effect of local costs from sorting on unobservables.

This interpretation motivates the model developed in Section 4. Rather than treating each correlation in isolation, we build a framework in which households make joint decisions over location,

commuting, housing, educational investment, and fertility, with local wages and housing prices clearing in equilibrium. The model allows us to trace observed fertility patterns to their structural sources and to evaluate policies that operate through different channels. A housing subsidy in Seoul, for example, may attract additional high-skilled workers to the capital region, intensifying education competition and partially offsetting the fertility gains from lower housing costs. This equilibrium logic is central to the policy analysis in Section 7.

4 Model

We consider an overlapping generations (OLG) framework where the population is comprised of three distinct stages of life: children (c), young adults (y), and the old (o). The model captures the lifecycle of a household through a sequential decision-making process. First, young adults, acting as parents, choose their residential neighborhood and commuting zone while simultaneously determining their desired fertility. Second, conditional on the number of children, parents make educational investment decisions. Finally, the labor and housing markets clear based on these aggregated individual choices and the existing population of the old.

4.1 Young Adults

The residential and labor market decisions of young adults are driven by an indirect utility function. For a household with education level e and n' children, the indirect utility derived from living in neighborhood j and working in commuting zone k is defined as:

$$U_{j,k}^{e,n'} = A_j^{e,n'} - C_{j,k} + \log\left((1-\tau_p)Y_k^e\right) - \alpha^{e,n'} \log p_j \tag{1}$$

In this formulation, $A_j^{e,n'}$ represents the exogenous amenity value, which is specific to the parents' education and the number of children.[7] The term $C_{j,k}$ captures the physical and temporal costs of commuting between the residence and the workplace. Income is represented by Y_k^e, reflecting the earnings of a worker with education e in zone k, and a τ_p fraction of income is saved into pension funds. Lastly, $\alpha^{e,n'} \log p_j$ denotes the disutility from housing expenditures, where p_j is the local housing price, and $\alpha^{e,n'}$ the housing expenditure share.

[7] We can endogenize the amenities by introducing agglomeration forces.

The lifetime value of a young adult is determined by the optimal selection of a neighborhood j' and a commuting zone k'. We assume households face an idiosyncratic taste shock $\varepsilon'_{j,k}$. The value function is given by:

$$V^y(e,j) = \mathbb{E}_{\varepsilon'_{j,k}}\left[\max_{j',k'} \sigma_{JK}\varepsilon'_{j,k} - T^e_{j,j'} + V^{y,c}(e,j',k') + \beta^o V^o(e,j')\right] \tag{2}$$

Here, $T^e_{j,j'}$ represents the cost of moving from the childhood neighborhood j to the adult neighborhood j'. The parameter σ_{JK} scales the taste shocks. The agent's total value is a discounted sum of the expected value of their children including their own contemporaneous utility, $V^{y,c}$, and the value of their own old age, V^o, with a discount factor β^o.

The fertility decision is nested within the household's problem. When choosing the number of children $n' \in \{0, 1, \ldots, N\}$, the household maximizes over the sum of current utility and the future prospects of their offspring:

$$V^{y,c}(e,j',k') = \mathbb{E}_{\varepsilon'_n}\left[\max_{n'}\left\langle \sigma_N\varepsilon'_n + U^{e,n'}_{j',k'} + \mathbb{I}(n' \geq 1)\{W^c(e,j',n') - F^{e,n'}_{j'}\}\right\rangle\right] \tag{3}$$

In this context, $F^{e,n'}_{j'}$ is a fixed cost associated with child-rearing, and $W^c(e,j',n')$ is the aggregate expected value of all n' children after their education level has been determined. This aggregate value is calculated as $W^c(e,j',n') = \sum_{m=1}^{n'} W^c_m(e,j',n')$ where $W^c_m(e,j',n')$ is the value of each child m.

Parents invest in their children's education level e' to maximize the child's future value $V^y(e',j')$, subject to an average education cost per child Φ, which varies across different neighborhood j', parents' education e, children's education e', and the number of children n'. The value of the n'-th child is expressed as:

$$W^c_m(e,j') = \mathbb{E}_{\varepsilon_{e'}}\left[\max_{e'} \sigma_E\varepsilon_{e'} - \Phi^{e,e',n'}_{j'} + \beta^y V^y(e',j')\right] \tag{4}$$

Parents choose the optimal level of education investment $e' \in \{L, H\}$ for their children, considering the education costs and associated future benefits of their children $V^y(e',j')$. Parents will make the educational investment ($e' = H$) as long as children can earn higher wages when they

become workers in neighborhood j'.

4.2 The Old Generation

For the final stage of life, the value of being old depends on neighborhood amenities and the level of pension benefits b_j^e, adjusted for housing costs:

$$V^o(e,j) = A_j^{e,n} + \log b_j^e - \alpha^{e,n} \log p_j \tag{5}$$

The pension benefits b_j^e are distributed according to the defined income replacement rate. We provide a detailed explanation of the pension system in the following section.

4.3 Decision Rules

The model is characterized by a sequence of discrete choices made by households, where idiosyncratic preferences follow a Type-1 Extreme Value distribution. Each household i draws independent and identically distributed idiosyncratic taste shocks for child preferences ($\{\epsilon_n^i\}_n$), location choice ($\{\epsilon_{j,k}^i\}_{j,k}$) and education levels ($\{\epsilon_e^i\}_e$) from a Type-I extreme value distribution.[8] This leads to the familiar multinomial logit structure for location, fertility, and education choices.

Location Choice and Commuting Patterns. The probability that a household with education level e, originating from neighborhood j, selects neighborhood j' for residence and commuting zone k' for work is defined as:

$$\lambda_{j,j',k'}^e = \frac{\exp\left(-T_{j,j'}^e + V^{y,c}(e,j',k') + \beta^o V^o(e,j')\right)^{\frac{1}{\sigma_{JK}}}}{\sum_{j'',k''} \exp\left(-T_{j,j''}^e + V^{y,c}(e,j'',k'') + \beta^o V^o(e,j'')\right)^{\frac{1}{\sigma_{JK}}}} \tag{6}$$

where σ_{JK} is the scale parameter governing the dispersion of location-job taste shocks. A key analytical result of this structure is the determination of commuting odds. By comparing two potential commuting zones k'' and k' for a household in a fixed neighborhood j', we can isolate the

[8]The probability density function is defined as $F(x) = \exp(\exp(-x - \bar{\gamma}))$ where $\bar{\gamma}$ denotes the Euler-Mascheroni constant. The parameters σ_N, σ_{JK} and σ_E serve to flexibly parameterize the effective variance by pre-multiplying these shocks across all model equations. This specific functional form is adopted because it allows for the derivation of tractable analytical expressions for both the value functions and the corresponding population shares of agents making each discrete choice.

labor market returns:

$$\underbrace{\log \frac{\lambda^e_{j,j',k''}}{\lambda^e_{j,j',k'}}}_{\text{Commuting Odds}} = \underbrace{\frac{1}{\sigma_{JK}}\left(V^{y,c}(e,j',k'') - V^{y,c}(e,j',k')\right)}_{\text{Return from job relocation}} \tag{7}$$
$$= \frac{1}{\sigma_{JK}}(\log \frac{Y^e_{k''}}{Y^e_{k'}} - (C_{j,k''} - C_{j,k'}))$$

See Appendix D.1 for derivation. This return from job relocation illustrates that commuting patterns are driven purely by the differential value of young adulthood across zones, as residential amenities and moving costs cancel out in the log-odds ratio. Intuitively, people can earn higher wages or reduce commuting costs by relocating their commuting zone, which is the key motive for reallocation.

Fertility Choice and the Value of Children. Depending on their spatial choices, households determine their optimal number of children $n' \in \{0, 1, \ldots, N\}$. The probability of having n' children is as follows.[9]

$$\gamma^{e,n'}_{j',k'} = \frac{\exp(U^{e,n'}_{j',k'} + \mathbb{I}(n' \geq 1)\{W^c(e,j',n') - F^{e,n'}_{j'}\})^{\frac{1}{\sigma_N}}}{\sum_{n=0}^{N} \exp(U^{e,n}_{j',k'} + \mathbb{I}(n' \geq 1)\{W^c(e,j',n) - F^{e,n'}_{j'}\})^{\frac{1}{\sigma_N}}} \tag{8}$$

From this, the model predicts the log-odds of having one child ($n = 1$) relative to remaining childless ($n = 0$) as follows:

$$\log\left(\frac{\gamma^{e,1}_{j,k}}{\gamma^{e,0}_{j,k}}\right) = \frac{1}{\sigma_N}\left(A^{e,1}_j - A^{e,0}_j + W^c(e,j,1) - F^{e,1}_j\right) \tag{9}$$

This log-odds ratio indicates that the propensity to have a child increases as the net benefit-defined as the amenity gain ($A^{e,1}_j - A^{e,0}_j$) plus the child's value $W^c(e,j,1)$ minus the fixed cost $F^{e,1}_j$increases. Notably, once the residential neighborhood (j) and workplace (k) are controlled for, the fertility decision becomes independent of income and housing prices. Instead, it is driven solely by idiosyncratic amenity gains, the value of the child, and fixed child-rearing costs.

[9]The associated inclusive value, which represents the expected value of young adulthood $V^{y,c}$, is the log-sum of these options: $V^{y,c}(e,j',k') = \sigma_N \log\left[\sum_{n=0}^{N} \exp(U^{e,n}_{j',k'} + \mathbb{I}(n' \geq 1)\{W^c(e,j',n) - F^{e,n}_{j'}\})^{\frac{1}{\sigma_N}}\right]$.

Education Choice. In the final stage of the households decision-making process, parents determine the education level e' for each of their n' children. This choice is modeled as a discrete investment problem where parents maximize the child's expected lifetime value, net of the costs associated with educational attainment. We assume that childrens outcomes are influenced by an idiosyncratic shock $\varepsilon_{e'}$, representing child-specific aptitude or preference, which follows a Type-1 Extreme Value distribution with a dispersion parameter σ_E. For each n'th child, the resulting probability $\pi_{j'}^{e,e',n'}$ that a parent with education level e residing in neighborhood j' will have a child who achieves education level e' is derived by the following multinomial logit structure:

$$\pi_{j'}^{e,e',n'} = \frac{\exp\left(-\Phi_{j'}^{e,e',n'} + \beta^y V^y(e', j')\right)^{\frac{1}{\sigma_E}}}{\sum_{e''} \exp\left(-\Phi_{j'}^{e,e'',n'} + \beta^y V^y(e'', j')\right)^{\frac{1}{\sigma_E}}} \tag{10}$$

This specification highlights several critical economic channels. The term $\Phi_{j'}^{e,e',n'}$ captures the neighborhood-specific cost of acquiring education e' for parents with education e for each of their n' children. By allowing this cost to depend on the parent's own education level e, the model accounts for intergenerational persistence; children of more educated parents may face lower effective costs due to home-produced human capital or superior information. Furthermore, the inclusion of n' in the cost function introduces a resource dilution effect, often referred to as the quantity-quality tradeoff, where a larger number of siblings increases the marginal cost of educational investment per child.

The forward-looking behavior of the household is encapsulated in the term $\beta^y V^y(e', j')$. Parents evaluate the return on education based on the childs future value as a young adult. Since $V^y(e', j')$ represents the value function of a young adult with education e' starting in neighborhood j', it intrinsically incorporates expected future wages, amenities, and commuting costs the child will encounter. Finally, the parameter σ_E governs the elasticity of education choice. A lower σ_E implies that educational attainment is highly sensitive to economic incentives, such as local labor market premiums, while a higher σ_E suggests that outcomes are driven more by idiosyncratic talent and shocks, effectively serving as a measure of social mobility within the model.

4.4 Population Dynamics

The aggregate population then evolves according to intergenerational laws of motion. The young population L^y at time t is composed of individuals who were children in the previous period:

$$L^{y,e}_{j',t} = \sum_{j} \sum_{k'} \lambda^{e}_{j,j',k'} L^{y,e}_{j,t} = \sum_{j} \sum_{k'} \lambda^{e}_{j,j',k'} L^{c,e}_{j,t-1} \tag{11}$$

The child population L^c is generated via fertility choices γ and the education transition π:

$$L^{c,e'}_{j',t} = \sum_{e} \sum_{k'} \sum_{n'} \frac{\gamma^{e,n'}_{j',k'} \cdot n'}{2} \pi^{e,e'}_{j'} L^{y,e}_{j',k',t} \tag{12}$$

The unit of analysis is the household, consisting of a man and a woman. To maintain a stationary population, each household needs to have two newborn children and there is no mortality between the young adult and old age periods. We abstract from endogenous matching (or marriage); households are formed exogenously. This assumption captures the strong degree of assortative matching in Korea; as shown in Table A.1, 76 percent of couples share the same educational attainment (i.e., both are either college graduates or non-college graduates).

4.5 Pension System

The pension system is modeled as a balanced budget framework that equates the total accumulated contributions from the young generation to the aggregate benefits distributed to the old. We define τ_p as the mandatory contribution rate and b^e_j as the pension benefit designated for education group e, which is typically calibrated as a fixed replacement rate of the worker's young-age wage. To ensure fiscal solvency in a stationary equilibrium, the system must satisfy the following budget condition:

$$\underbrace{\left(\sum_{e,k} \tau_p \cdot Y^e_k \cdot L^{y,e}_k \right) (1+R)}_{\text{Total Contributions with Return}} = \underbrace{\sum_{e,j} b^e_j \cdot L^{o,e}_j}_{\text{Total Benefits}} \tag{13}$$

Here, benefits per household b^e_j are defined as a fraction ρ_p of their pre-retirement income. In this formulation, the left-hand side represents the total fiscal inflow, where Y^e_k is the gross labor income of young agents of education level e in commuting zone k, and $L^{y,e}_k$ is their corresponding

population mass. The parameter R represents the return on the pension fund; in the context of a steady-state Pay-As-You-Go (PAYG) system, R is the required return of the pension fund, which is endogenously determined by the pension fund budget balance. The right-hand side represents the total expenditure paid to the old population $L_j^{o,e}$.

4.6 Local Production

Production in commuting zone k follows a Constant Elasticity of Substitution (CES) technology that combines low-skilled (l) and high-skilled (h) labor:

$$Q_k = Z_k \left(S_k^l (L_k^l)^{\frac{\rho-1}{\rho}} + S_k^h (L_k^h)^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}} \tag{14}$$

Assuming competitive markets, the income Y_k^e for a worker of education type e is determined by the marginal product of labor:

$$Y_k^e = Z_k \cdot S_k^e \cdot (L_k^e)^{-\frac{1}{\rho}} \cdot \left(S_k^l (L_k^l)^{\frac{\rho-1}{\rho}} + S_k^h (L_k^h)^{\frac{\rho-1}{\rho}} \right)^{\frac{1}{\rho-1}} \tag{15}$$

where Z_k is total factor productivity, S_k^e represents education-specific productivity shifters, and ρ is the elasticity of substitution between skill types ($e = \ell$ or h).

4.7 Equilibrium

To close the model, we define the requirements for a stationary competitive equilibrium. In this state, all agents optimize their objectives given local prices, and the aggregate distributions of population remain invariant.

Definition 1 (Stationary Competitive Equilibrium)**.** A stationary competitive equilibrium is defined as a set of population distributions $\{L_j^{c,e}, L_j^{y,e}, L_j^{o,e}\}$, endogenous income levels $\{Y_k^e\}$, location and cost parameters $\{A_j^{en}, C_{j,k}, F_j^{e,n}, \Phi_{j'}^{e,e',n'}, \bar{H}_j^s\}$, and local prices for wages $\{w_k\}$ and housing $\{p_j\}$, such that:

- **Household Optimization:** Given education e, initial neighborhood j, and prices $\{w_k, p_j\}$, households choose residence j', commuting zone k', fertility n', and children's education e' to maximize lifetime utility according to (2), (3), and (4).

- **Firm Optimization:** Firms in each commuting zone k maximize profits by choosing labor demand $L_k^{d,e}$ for each education level e such that the wage w_k equals the marginal product of labor defined in (15).

- **Population Distribution:** The aggregate dynamics is consistent with the migration decisions of individuals (6). Specifically, the distribution of young adults $L_j^{y,e}$ and their children $L_j^{c,e}$ across the neighborhoods j is stationary, given the old $L_j^{o,e}$.

- **Pension Fund:** Given the contribution rate τ_p and the income replacement rate ρ_p, the pension fund satisfies the intergenerational budget constraint (13) by endogenously determining the fund return R.

- **Market Clearing:**

    1. **Labor Market:** In each zone k, the sum of labor supplied by young adults matches firm demand: $\sum_j L_{j,k}^{y,e}(w_k) = L_k^{d,e}(w_k)$.
    2. **Housing Market:** In each neighborhood j, the aggregate housing demand from young and old households equals the fixed supply: $\sum_e \sum_k [H_{j,k}^{y,e}(p_j) + H_{j,k}^{o,e}(p_j)] = \bar{H}_j^s$.

The definition of the stationary equilibrium presented above assumes a fixed distribution of the elderly population, abstracting from transitional dynamics. We employ this equilibrium concept to calibrate the model in Section 5 and to conduct counterfactual policy analyses by comparing alternative steady states in Section 7. However, we find that a long-run steady state exists where the elderly population size is determined endogenously according to the intergenerational laws of motion, (11) and (12). Consequently, we also perform transition path analyses during this intergenerational adjustment in Section 7.6.

5 Parameterization

The structural parameters of the model are categorized into those estimated directly from microdata and those calibrated based on the existing literature or standard economic benchmarks.

5.1 Parameter Classification and Identification

Table 2 details the cost and preference parameters estimated based on the model's theoretical predictions. To be specific, the identification strategy relies on the multinomial logit structure of the agents' choices. For instance, moving costs ($T^e_{jj'}$) are identified by the sensitivity of migration flows to bilateral distance, while the variance of taste shocks (σ_{JK}, σ_N) is recovered by regressing choice probabilities on their respective economic returns (income differentials and amenity gains).

Table 2 also outlines the calibrated parameters. Given the estimated parameters, the altruism parameter β^y is set to match the empirical average fertility rate. The parameter for education taste shocks (σ_E) is also simultaneously calibrated to match observed unconditional dispersion in education outcomes ($\pi^{e,e',n'}_{j'}$). The long-run discount factor β^o is set to 0.290, reflecting a 30-year generational gap with an annual discount rate of 0.96. Technology parameters, such as the elasticity of substitution (ρ), are drawn from Acemoglu and Autor (2011). Following Eckert and Kleineberg (2024), local productivity (Z_k, S^h_k) and housing supply ($\bar{H}^s_j$) are backed out as residuals from the firms' first-order conditions and housing market-clearing identities, respectively.

Table 2: Parameterization

Symbol	Value	Description	Source / Estimation Method
I. Estimated Cost and Preference Parameters			
$\Phi^{e,e',n'}_j$	—	Education Costs	Regression (spending/income proxy)
$T^e_{jj'}$	—	Moving Costs	Regression (distance/province)
$C_{j',k'}$	—	Commuting Costs	Regression (distance proxy)
$A^{e,n'}_j$	—	Local Amenities	Neighborhood FE (birth regression)
$F^{e,n'}_j$	—	Child-rearing Fixed Costs	Neighborhood FE (birth regression)
II. Taste Shocks (Unobserved Heterogeneity)			
σ_{JK}	0.12	Mobility Shock Std. Dev.	IV regression (commuting log-odds)
σ_N	0.75	Fertility Shock Std. Dev.	Log-odds regression (birth choice)
σ_E	0.28	Education Shock Std. Dev.	Target: Education attainment dispersion
III. Discount and Altruism Factors			
β^y	0.004	Intergenerational Altruism	Target: Mean fertility rate $\mathbb{E}[\gamma]$
β^o	0.290	Future Value Discount	Generational benchmark (30-yr horizon)
IV. Production and Housing			
ρ	1.50	Substitution Elasticity	Acemoglu and Autor (2011)
Z_k, S^h_k	—	Spatial Productivity	Labor demand FOC (wage-population data)
$\bar{H}^s_j$	—	Housing Supply	Housing demand FOC (housing prices)
$\alpha^{e,n}$	—	Housing Budget Share	Observed expenditure data

5.2 Estimation of Education Costs

The structural parameters governing human capital investment are identified by exploiting the discrete choice problem faced by parents. We assume that households determine their childrens education level e' by maximizing the expected future value of that education, net of idiosyncratic shocks and neighborhood-specific costs.

Econometric Strategy and Measurement. To recover the parameters of the education cost function, we employ a Poisson Pseudo-Maximum Likelihood (PPML) regression following Silva and Tenreyro (2006). This estimator is robust to the presence of zero-frequency cells in the data and accounts for the non-linear relationship between costs and choice probabilities. After rearranging the decision rule for education, the empirical specification is defined as:

$$\pi_{\mathcal{J}}^{e,e',n'} = \exp(\delta_{\mathcal{J}}^{e'} + \psi_{\mathcal{J}}^{e} - \iota^{e,n'} X_{\mathcal{J}}^{e,e'}) + \tilde{\varepsilon}_{\mathcal{J}}^{e,e'} \tag{16}$$

where $\delta_{\mathcal{J}}^{e'}$ and $\psi_{\mathcal{J}}^{e}$ are province-level fixed effects for child and parent education, respectively. The term $\delta_{\mathcal{J}}^{e'}$ absorbs the discounted future value $\beta^y v^y(e', j')$, while $\psi_{\mathcal{J}}^{e}$ absorbs the inclusive value term in the denominator.[10] The cost proxy $X_{\mathcal{J}}^{e,e'}$ is measured using the annual Private Education Survey (2020–2024), defined as the ratio of annual education spending per child to annual household income. See Appendix D for details.

Our identification strategy rests on the assumption that education costs are uniform within a province $\mathcal{J}$ (i.e., *si-do*), even as residential choices vary at the more granular neighborhood j (*dong*) level. Specifically, we assume $\phi_{j'}^{e,e',n'} = \phi_{j''}^{e,e',n'}$ for all $j', j'' \in \mathcal{J}$, which allows us to recover the estimated cost function as $\hat{\phi}_{\mathcal{J}}^{e,e',n'} = \hat{\iota}^{e,n'} X_{\mathcal{J}}^{e,e'}$. In the following section, we estimate the key structural parameters–including moving costs and amenities–at the provincial level.

Results and Interpretation. The results of the PPML regression, presented in Table 3, reveal a significant negative relationship between education costs and attainment. For college-educated parents, a 1 percentage point increase in the education cost share leads to a 5.60% decrease in the odds of the child achieving a high-education outcome.

[10] We assume that parent education-province fixed effects effectively control for the denominator. Given our sample size, we avoid education-province-children ($e-\mathcal{J}-n'$) fixed effects, as these would absorb the majority of the variation

The interaction terms demonstrate that this deterrent effect is locally convex in the number of children, providing structural evidence of a resource dilution effect within the household. In high-cost regions like Seoul, these costs contribute to an "educational rat race," where the budget share required to ensure competitive educational outcomes for children exerts substantial downward pressure on both the probability of attaining high-type education and the overall fertility rate.

Our empirical analysis reveals significant regional and educational heterogeneity in the costs associated with child-rearing. First, we find that education costs per child are locally convex in the number of children in Figure 2. This pattern is consistent with the "educational rat race" hypothesis proposed by Kim et al. (2024), where parents face increasing competitive pressures to invest in their children's human capital. In addition to their findings, we further demonstrate that the convexity of education costs is particularly pronounced in the metropolitan area of Seoul. This pattern serves as an important ingredient in the model, shaping the observed regional fertility gradients. Figure 3 validates our structural education cost estimates against empirical data. We find that the estimated costs are closely aligned with observed education spending as a share of household income. Notably, the model-implied costs are slightly higher than the raw spending data, as our estimates also capture the non-pecuniary costssuch as time costs and psychological pressureassociated with intensive education investments in South Korea.

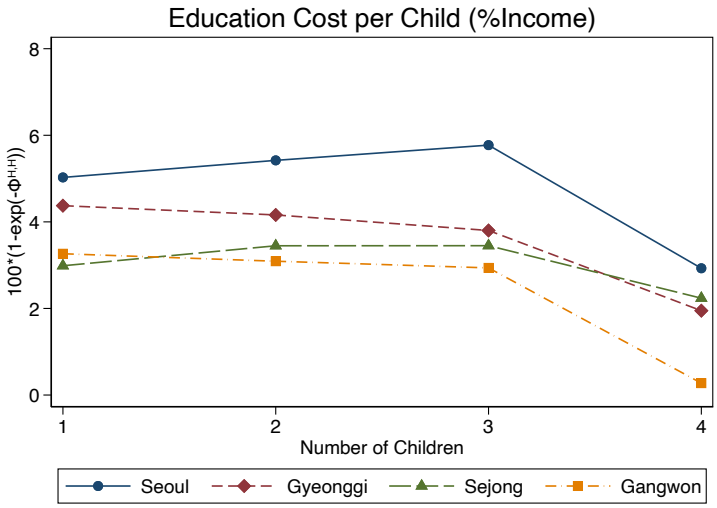
in education costs necessary for identification.

Table 3: Impact of Education Costs on Educational Outcomes ($\pi_{\mathcal{J}}^{e,e',n'}$)

	(1) College-educated	(2) Non-college-educated
Edu Cost ($X_{\mathcal{J}}^{e,e',n'}$)	-5.60*** (0.57)	-6.26*** (1.08)
$X \times \mathbb{I}[n'=2]$	-0.81*** (0.14)	-1.12*** (0.09)
$X \times \mathbb{I}[n'=3]$	-1.50*** (0.19)	-2.84*** (0.27)
$X \times \mathbb{I}[n'=4]$	0.20 (0.19)	-1.50*** (0.57)
Province FE	Yes	Yes
Child Edu × Province FE	Yes	Yes
Observations	2,011	2,124
Pseudo R^2	0.0438	0.0570

Notes: Standard errors in parentheses, clustered at province and year. *** p<0.01.

Figure 2: Education Costs by Region



Notes: Figure shows that education costs reach their peak in Seoul, particularly for households with college-educated parents. We express these costs as the percentage of income required to compensate for the utility or outcome loss associated with education expenditures, effectively capturing the shadow price of child quality.

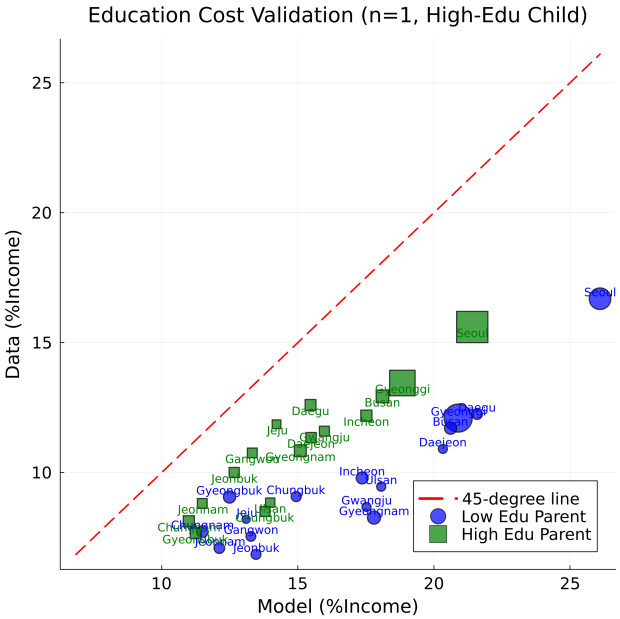


Figure 3: Education Costs: Model vs. Data

Notes: Figure plots education spending per child (relative to income) observed from data. Model implied education costs per child (when households have one child, $n' = 1$) are converted into the fraction of income.

5.3 Estimation of Moving Costs

To quantify the barriers to migration, we estimate bilateral moving costs, $t^e_{jj'}$, using a gravity-based framework applied to Korean population census in 2020. Following the structural implications of the model, the probability $\lambda^e_{j,j',k'}$ that an individual of education type $e \in \{H, L\}$–originally located in neighborhood j–migrates to j' and commutes to k' is given by:

$$\lambda^e_{j,j',k'} = \frac{\exp(-t^e_{jj'}) \exp\left(v^{y,c}(e, j', k') + \beta^o v^o(e, j')\right)}{\sum_{m,n} \exp(-t^e_{jm} + v^{y,c}(e, m, n) + \beta^o v^o(e, m))} \tag{17}$$

where the normalized moving cost is $t^e_{jj'} (\equiv T^e_{jj'}/\sigma_{JK})$. The remaining terms in the numerator represent the expected lifetime value of residing in j' and working in k', scaled by the dispersion of location taste shocks σ_{JK}. Following Eckert and Kleineberg (2024) and the standard spatial literature, we recover these costs by estimating a Poisson Pseudo-Maximum Likelihood (PPML) regression for each parental education e:

$$\lambda^e_{j,j',k'} = \exp\left(\delta^e_j + \psi^e_{j',k'} - \iota^e_1 \text{dist}_{jj'} - \iota^e_2 \mathbb{I}(\text{prov}_j \neq \text{prov}_{j'})\right) + \varepsilon^e_{j,j',k'} \tag{18}$$

In this specification, δ^e_j denotes origin fixed effects, while $\psi^e_{j',k'}$ captures destination-commuting zone fixed effects. The moving cost $t^e_{jj'}$ is predicted as a function of the bilateral distance between locations and a dummy variable for moves within the same province.

Empirical Results. Table 4 presents the PPML estimation results. Consistent with standard gravity equations, distance significantly increases the moving cost. The estimates suggest that a 100km increase in distance between two regions leads to a drop in the migration share of 1.8 percentage points for high-skilled workers and 2.3 pp for low-skilled workers. Simply moving to a different province results in 5.0 pp and 5.4 pp drops in the migration share for each respective group. Notably, we find that individuals with lower educational attainment (non-college) face higher moving costs, suggesting greater spatial frictions for this demographic.

Table 4: PPML Estimation of Moving Costs

	(1) High-Skilled	(2) Low-Skilled
$\text{Distance}_{j,j'}$ (100km)	-0.68***	-0.84***
	(0.07)	(0.06)
$\mathbb{I}(\text{prov}_j \neq \text{prov}_{j'})$	-1.90***	-1.99***
	(0.11)	(0.08)
Marginal Effects on Migration Share		
Distance (100km)	-1.8 *pp*	-2.3 *pp*
Inter-province Effect	-5.0 *pp*	-5.4 *pp*
Observations	643	631

Notes: Neighborhood ($\mathcal{J}$) $\times$ commuting zone and origin fixed effects controlled. Destination clustered standard errors are in parentheses. Marginal effects represent the percentage point change in the dependent variable for a one-unit increase in the regressor.

5.4 Estimation of Commuting Cost

We estimate commuting costs in two steps. The first step recovers the dispersion parameter σ_{JK} from an instrumental variables regression on the commuting log-odds ratio:

$$\log \frac{\lambda^e_{j,j',k''}}{\lambda^e_{j,j',k'}} = \frac{1}{\sigma_{JK}} \left(\log \frac{Y^e_{k''}}{Y^e_{k'}} - (C_{j',k''} - C_{j',k'}) \right). \tag{19}$$

OLS estimation of this equation is biased because local wages and commuting flows are jointly determined through labor market clearing. To address this, we construct a Bartik instrument from 2010 industry employment shares interacted with leave-one-out national industry wage growth across 77 two-digit industries, following Bartik (1991) and Borusyak et al. (2022). Taking first differences eliminates time-invariant location components, and the instrument isolates the component of local wage growth driven by national industry trends. Appendix D.2 provides full construction details and identification assumptions.

Table 5 reports the 2SLS estimates. The coefficient on the wage growth differential is $1/\hat{\sigma}_{JK} = 8.14$, implying $\hat{\sigma}_{JK} = 0.12$. Standard errors are clustered at the origin province level. The first-stage Kleibergen-Paap F-statistic is 32.3, well above the Stock and Yogo (2005) critical value of 16.38 for 10 percent maximal IV size. The Anderson-Rubin 95 percent confidence interval, which is robust to weak instruments, is $[0.4, 15.6]$. These diagnostics confirm the validity of our Bartik instrument.

In the second step, we assume commuting costs are linear in distance, $C_{j',k} = \xi \cdot \text{dist}_{j',k}$, and

Table 5: Bartik IV Estimation of Commuting Elasticity

	(1) $\Delta \log(\lambda^e_{k''}/\lambda^e_{k'})$
$\Delta(g^e_{k''} - g^e_{k'})$	8.136**
	(3.524)
Kleibergen-Paap F-stat	32.3
Observations	1,749

Notes: 2SLS estimates. The dependent variable is the time-differenced log commuting odds ratio. The endogenous regressor is the difference in log wage growth across commuting zone pairs. The instrument is the corresponding difference in Bartik predicted wage growth, constructed from 2010 industry employment shares and leave-one-out national industry wage growth across 77 two-digit industries. Standard errors clustered at the origin province level (17 clusters) in parentheses. The sample pools 2010–2015 and 2015–2020 differences across 15 commuting zones defined at the province (*sido*) level, restricted to canonical pairs with non-missing commuting flows. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

recover ξ by regressing the residualized log commuting ratio on the distance differential by OLS (Table A.2 in Appendix D.3).

5.5 Estimation Strategy for Amenities and Fixed Costs

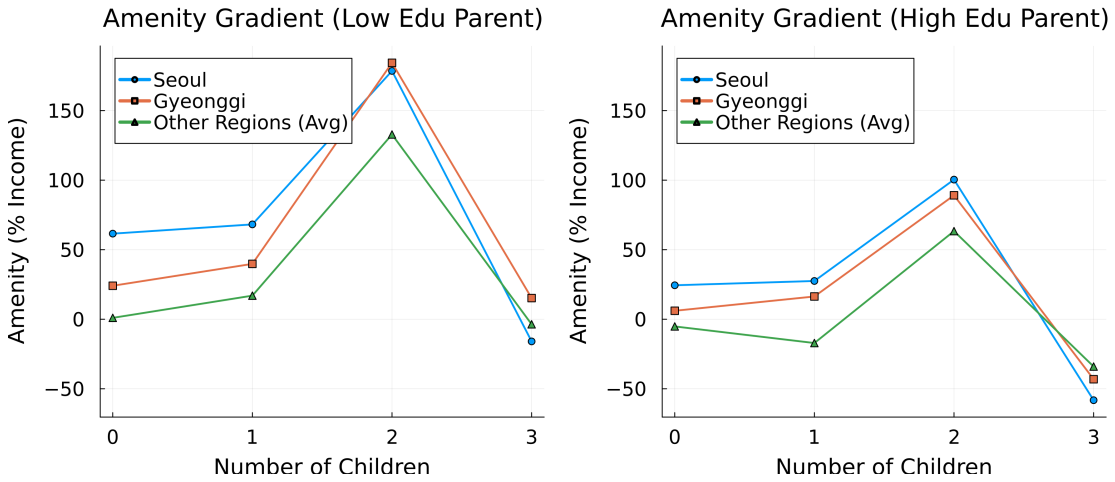
The estimation of local amenities (A) and fixed costs of children (F) exploits the birth rate data ($\gamma^{e,n'}_{j',k'}$) from the Census. We normalize the indirect utility (U), amenities and the fixed costs by the taste shock parameter: $u \equiv U/\sigma_N$, $a \equiv A/\sigma_N$, $f \equiv F/\sigma_N$, and $c \equiv C/\sigma_N$.

For households with children ($n' > 0$), the log birth rate incorporates the value of children (w^c) and the fixed costs of child-rearing:

$$\log \gamma^{e,n'}_{j',k'} = \underbrace{[a^{e,n'}_{j'} - f^{e,n'}_{j'} + \frac{\sigma_E}{\sigma_N} w^c(e,j',n')]}_{\text{Edu}\times\text{Neighborhood}\times\text{Number of children FE}} + \frac{1}{\sigma_N}[\log(\frac{Y^e_k}{p^{\alpha^{e,n'}}_{j'}}) - \sigma_{JK} v^{y,c}(e,j',k') - C_{j'k'}] \tag{20}$$

Given the fertility choice n', we recover the composite of destination amenities $a^{e,n'}_{j'}$, fixed costs $f^{e,n'}_{j'}$, and the value of children w^c by estimating education-neighborhood-children fixed effects ($e \times j' \times n'$). Once w^c is identified through this procedure, we are able to separately calculate the remaining combination of amenities and fixed costs. See Appendix E for details. The regression yields an estimate for the taste shock parameter $\hat{\sigma}_N = 1/0.793 \approx 1.26$ where the denominator is estimated as the regression coefficient from Table A.3.

Figure 4: Estimated amenities



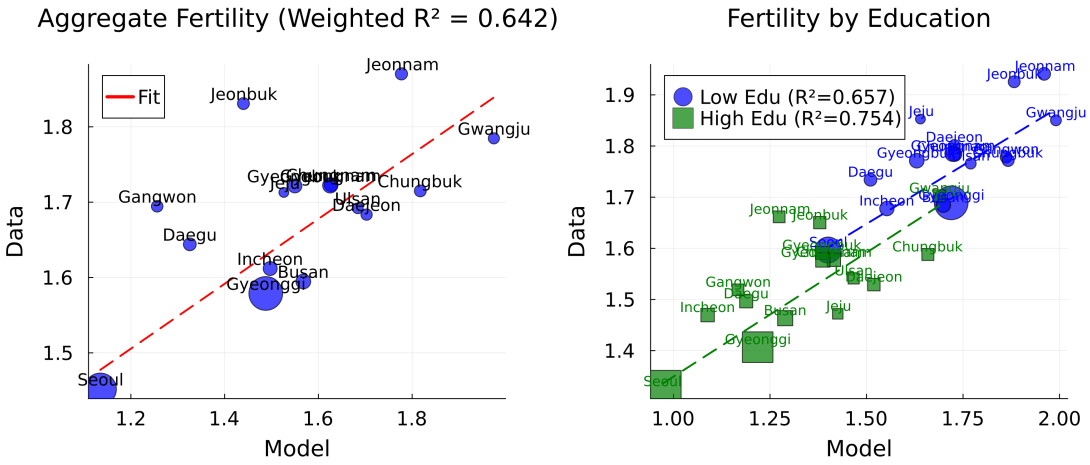
Notes: The figure displays the estimated net amenities (amenities minus fixed costs) for low- and high-education households by province. We express these values as the percentage of income required to compensate for the utility change associated with the net amenities in each location.

Our empirical results reveal significant spatial and demographic patterns. We find that net amenity gains increase only up to the second child ($n' = 2$), after which fixed costs dominate the gains for larger families. Net amenity gains are found to be relatively uniform across parental education levels. Furthermore, urban areas like Seoul exhibit a higher value of children (w^c), likely driven by higher relative wages in these areas despite higher education costs. This suggests that the primary drivers of low fertility in Seoul are likely high education costs (ϕ) rather than a lower value for children.

5.6 Solution Method and Model Fit

Given the parameterization, we solve the model using a value function iteration (VFI) algorithm. Detailed descriptions of the solution method, including procedures for the counterfactual and transition analyses presented in the following sections, are provided in Appendix F. The model demonstrates a strong fit with regional data, particularly regarding fertility, where it effectively captures regional variations ($R^2 \approx 0.585$) and matches patterns across both low- and high-education groups in Figure 5. The model successfully replicates the extremely low fertility rate in the Metropolitan area: Seoul, Gyeonggi, and Incheon. Moreover, it reproduces lower fertility rates among high-skilled workers than those of low-skilled workers. While the fit for low-skill wages is reasonable, high-skill wages exhibit a lower correlation attributable to the model's mechanism of extreme sorting. Addi-

Figure 5: Model versus Data



Notes: We compare observed data and model-simulated data in the stationary equilibrium across regions.

tionally, although the model accurately reproduces the spike in housing prices in Seoul, it predicts somewhat higher concentration effects than those observed in the empirical data.

6 Decomposition: Decline in Fertility Rate

Methodology. To quantify the drivers of fertility change, we employ a structural decomposition analysis that partitions the total change in the national Total Fertility Rate (TFR) into three distinct components: a within-region effect, a compositional effect, and a covariance term. The within component captures changes in fertility behavior among households staying within their respective regions, reflecting their response to local price shocks such as rising education or housing costs. The compositional component measures the shift in the national average caused by the spatial sorting of households across regions with heterogeneous fertility baselines; for instance, if households migrate from high-fertility rural areas to low-fertility urban centers like Seoul, the national TFR falls even if regional fertility rates remain constant. Finally, the covariance term accounts for the simultaneous change in both regional population shares and regional fertility rates.

Decomposition. To understand why the national effect of a change in structural parameters or a policy can differ in sign from its local effect, we decompose changes in the national fertility rate.

$$\Delta \text{TFR} = \underbrace{\sum_j \Delta \text{TFR}_j \cdot s_{j,0}}_{\text{Within}} + \underbrace{\sum_j \text{TFR}_{j,0} \cdot \Delta s_j}_{\text{Composition}} + \underbrace{\sum_j \Delta \text{TFR}_j \cdot \Delta s_j}_{\text{Covariance}} \tag{21}$$

where $s_{j,0}$ is the baseline population share of region j, $\text{TFR}_{j,0}$ is the baseline fertility rate in region j, and Δ denotes the change from baseline to counterfactual. The within term captures behavioral fertility changes holding the population distribution fixed. The composition term captures changes in national fertility that arise purely from households moving between regions with different baseline fertility levels. The covariance term captures the interaction between fertility changes and population changes, and is typically small. The key insight from this decomposition is that a policy can raise fertility within the regions it targets (positive within effect) and still lower national fertility if it reallocates population toward regions where fertility is low (negative composition effect that dominates).

Table 6 presents the results of this decomposition for the 2005–2020 period, revealing that the sharp decline in South Korean fertility (from around 2.0 to 1.5) is primarily driven by within-region responses to deteriorating child-rearing amenities and rising education costs.[11] Deteriorating child-rearing amenities and rising education costs emerge as the dominant factors, contributing -0.080 and -0.043 points to the total decline, respectively, with their impact almost entirely concentrated in the *within* column. The preference shock (β^y) also play a substantial role, explaining -0.082 of the total change, which partly reflects the changes in gender roles or social norms over time. To sum up, while declining child-rearing amenities and increasing education costs are the primary factors—explaining roughly 23 percent of the decrease—a preference shock also plays a significant role, accounting for an additional 16 percent. Notably, while factors like income and housing supply show significant positive total effects, they are due to positive *compositional* effects, suggesting that the resulting spatial reallocation of the population provides a marginal, albeit insufficient, counterbalance to the overall downward trend. Figure 6 shows the estimated changes in child-rearing amenities and education costs between 2005 and 2020. Over this period, amenities for households

[11] We use completed cohort fertility as our primary measure, restricting the sample to households where the mother is aged 40 to 50.

with one child decreased significantly more than those for childless households. Moreover, education costs rose substantially during the same interval. These spatial and temporal patterns illustrate why deteriorating child-rearing amenities and rising education costs account for a substantial portion of the observed decline in fertility.

Furthermore, shifts in amenities favoring childless households ($n = 0$) account for a substantial 0.128-point decline in fertility. This trend potentially reflects an increasing societal preference for personal lifestyle over child-rearing. While improvements in housing supply and productivity had negligible impacts on fertility, the reduction in commuting costs—driven by transportation infrastructure enhancements—exerted a modest upward pressure on fertility by alleviating parental time constraints. Notably, demographic shifts in the population structure (the rising ratio of elderly to youth) reduced the fertility rate by an additional 0.126 points. The primary mechanism for this decline is observed in the housing market, where an aging population crowds out younger households, displacing them toward areas with inferior amenities.

Table 6: Decomposition of Fertility Change (2005–2020)

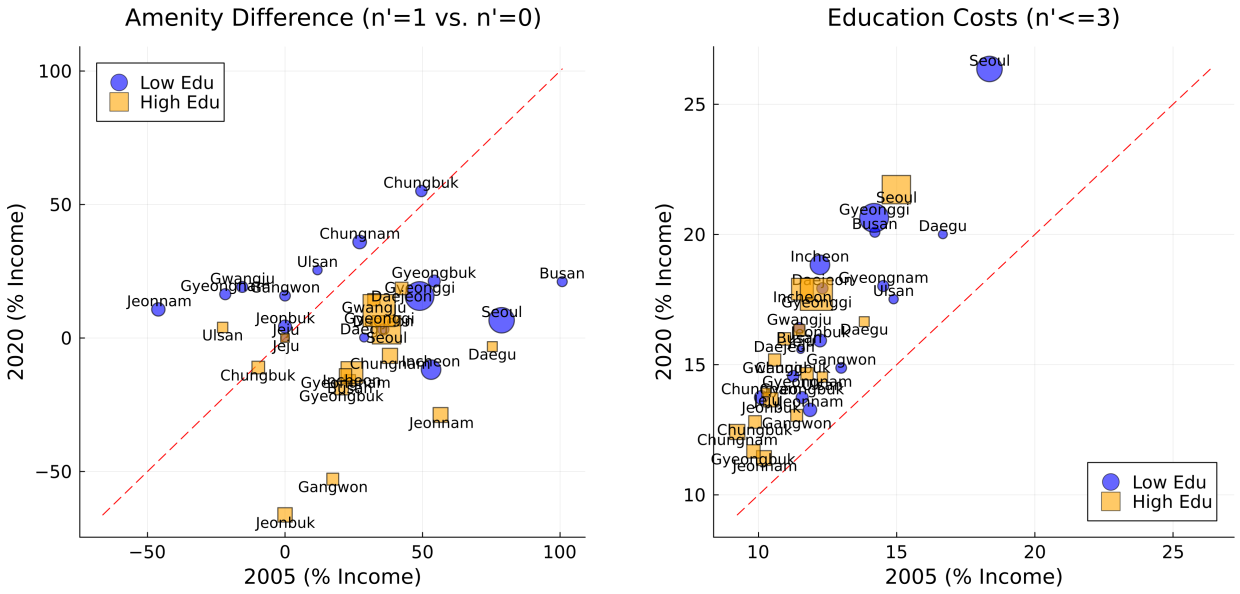
Factor	**Log %**	**Total**	**Within**	**Comp**	**Cov**
Altruism (β^y)	-4.097	-0.082	-0.074	-0.005	-0.002
Amenities (n≥1)	-4.011	-0.080	-0.110	+0.001	+0.029
Amenities (n=0)	-6.528	-0.128	-0.098	-0.015	-0.015
Education Costs	-2.145	-0.043	-0.052	+0.001	+0.008
Housing Supply	+0.088	+0.002	+0.002	-0.001	+0.001
Income (Productivity)	+0.749	+0.015	+0.001	+0.015	-0.000
Commuting Costs	-1.194	-0.024	-0.005	-0.022	+0.003
Population Structure	-6.413	-0.126	-0.077	-0.025	-0.024
Combined Effects					
Altruism + Amen.($n \geq 1$) + Inc + House + Comm	-6.541	-0.129	-0.177	+0.004	+0.044
Altru. + Amen.(all) + Inc + House + Comm	-14.264	-0.270	-0.286	-0.011	+0.027
Altru. + Amen.(all) + Inc + House + Comm + Pop	-26.007	-0.465	-0.447	-0.023	+0.006
Total Change (2005–2020)	-29.399	-0.517	-0.513	-0.015	+0.011

Notes: The table decomposes the observed change in the national TFR between 2005 and 2020. Components follow the decomposition identity: Total = Within + Composition + Covariance. Log % is $\ln(TFR_{2020}/TFR_{2005}) \times 100$.

7 Counterfactuals and Policy Experiments

We use the estimated model to evaluate fertility policies. The central question is not whether a given policy raises fertility in the location where it is implemented, but whether it raises fertility

Figure 6: Changes in Amenities and Education Costs



(a) Amenities ($A^{n'=1} - F^{n'=1} - A^{n'=0}$)

(b) Education Costs

Notes: Figure plots the values for amenities and education costs in 2005 and 2020.

nationally once households re-sort across locations and local prices adjust. This distinction matters because Korea's fertility problem is entangled with its spatial structure. The metropolitan area offers the highest wages but also the steepest education costs, highest housing prices, and longest commutes. Any policy that makes the metro area more attractive draws in households, which changes the local equilibrium that families face. Whether this re-sorting amplifies or offsets the direct policy effect is something only a general equilibrium framework can answer.

We organize the analysis around three sets of results. First, we examine place-based policies that target the Seoul metropolitan area and show that several commonly proposed interventions produce negligible or even negative effects on national fertility because the induced migration response offsets the within-region behavioral gains. Second, we evaluate nationwide education cost reduction and find that it dominates all other single-margin interventions. Third, we consider fiscal transfers and coordinated policy packages, showing that combining education reform with targeted subsidies generates the largest fertility response. Throughout, we decompose each policy's effect into within-region behavioral changes and cross-region compositional shifts, which reveals the mechanism behind each result.

7.1 Methodology and Decomposition

Counterfactual equilibrium. Each experiment modifies one or more structural primitives while holding preferences and migration frictions fixed. We then re-solve for a new stationary equilibrium in which wages and housing prices clear local markets and all young adults have made their migration, commuting, education investment, and fertility decisions according to the model's decision rules, given the existing old population. We report the national fertility rate and its log change relative to the baseline ($100 \times \ln(\text{TFR}_{cf}/\text{TFR}_{base})$) for each scenario. The fertility rate in the model is the total number of children born divided by the number of young adults, an average parity measure. This differs from the conventional period TFR (0.72 in 2023), which sums age-specific birth rates across single-year ages, so we report percentage changes throughout.

Shock implementation. We consider shocks to education costs, commuting costs, housing supply, amenities, and fiscal instruments, applied either to the Seoul metropolitan area alone or to all regions simultaneously. Housing supply shocks directly double the housing stock $\bar{H}_j^s$. Education cost, commuting cost, and amenity shocks are sized in consumption equivalence variation (CEV) units, meaning that a -20% CEV education cost shock reduces education costs by an amount equivalent in utility terms to a 20% reduction in consumption. This scaling makes shocks comparable across margins. Note that the education cost parameter Φ in the model captures the full cost of education competition, including time costs, opportunity costs, and the rat-race externality, not just direct spending on private tutoring. The effective welfare cost of education competition therefore substantially exceeds the 7–8% of household income that Seoul families spend directly on *hagwon*. Subsidy experiments modify the household budget constraint through taxes and transfers. For each place-based shock, we compare the Seoul-only version with the nationwide version to isolate the role of spatial reallocation.

Table 7 presents the main counterfactual results for the experiments discussed in this section.

7.2 Why place-based policies can backfire

A natural policy response to Seoul's low fertility is to make the metropolitan area more livable for families by expanding housing, shortening commutes, or improving childcare infrastructure. Before turning to the policies that work, we briefly illustrate why these commonly proposed interventions

Table 7: Counterfactual results: main policy experiments

	TFR	% Chg	Level Chg	Within	Comp	Cov
Baseline	1.493	–	–	–	–	–
Housing Supply (+100% Vol)						
Seoul only	1.466	-1.81%	-0.027	+0.011	-0.038	+0.000
All regions	1.499	+0.39%	+0.006	+0.006	+0.000	-0.000
Education Cost (−20% CEV)						
Seoul only	1.519	+1.72%	+0.026	+0.041	-0.079	+0.064
All regions	1.755	+16.16%	+0.262	+0.241	+0.021	+0.000
Commuting Cost (−20% CEV)						
Seoul only	1.453	-2.72%	-0.040	+0.009	-0.052	+0.003
All regions	1.494	+0.07%	+0.001	+0.001	+0.000	+0.000
Amenity Shock (+20% CEV)						
Seoul only	1.490	-0.21%	-0.003	+0.020	-0.035	+0.012
All regions	1.579	+5.58%	+0.086	+0.080	+0.006	-0.000
Subsidies to Parents						
Childless tax (20%)	1.630	+8.79%	+0.137	+0.129	+0.008	-0.000
Universal tax (rev. neutral)	1.536	+2.83%	+0.043	+0.040	+0.003	+0.000
Nationwide Edu Reform + Subsidies						
Edu + childless tax	1.855	+21.69%	+0.362	+0.336	+0.027	-0.001
Edu + universal tax	1.789	+18.08%	+0.296	+0.272	+0.023	+0.001

Notes: Each row reports the national fertility rate (TFR), its percentage change ($\ln(TFR_{cf}/TFR_{base}) \times 100$), and the decomposition of the level change into within, composition, and covariance components.

fail at the national level.

Every Seoul-only policy we consider raises fertility within the metropolitan area, but it also attracts households into a region where baseline fertility is low. This compositional shift offsets the behavioral gain, leaving national fertility unchanged or slightly lower. A commuting cost reduction of 20% (CEV) in Seoul lowers national fertility by 2.72% (0.040 points). The within effect is almost zero (0.009) but the composition effect (−0.052) dominates as lower commuting costs draw households into the metro area. Doubling Seoul's housing stock produces a negative effect (−1.81%). Housing prices fall by 30 percent in Seoul, inducing roughly 9 percentage points of population inflow, but the within gain (+0.011) is dominated by composition loss (−0.038). An amenity improvement of 20% (CEV) in Seoul for households with children lowers the fertility by 0.21 percent through the same mechanism.

The common thread is straightforward. Seoul is a low-fertility region in equilibrium because of intense education competition, high housing costs, and long commutes. Any policy that in-

creases Seoul's attractiveness without addressing these structural drivers pulls more households into a low-fertility environment. The within-region behavioral gain from the policy is real, but the compositional reweighting toward a low-fertility destination offsets it. A partial equilibrium evaluation, or a reduced-form study that estimates the local treatment effect without accounting for migration, would miss this offset entirely.

When the same policies are applied nationwide, the composition effect is muted because there is no single low-fertility destination pulling households in one direction. National commuting cost reduction raises fertility by 0.07%, national housing supply expansion by 0.39%, and national amenity improvements by 5.58%. All are positive but modest in the case of commuting costs and housing supply expansion. None approaches the magnitude of nationwide education cost reduction, which we turn to next.

7.3 Education cost reduction

Reducing education costs nationwide by 20% (CEV) raises the national fertility rate by 16%, the largest single-margin intervention in the model (Table 7).[12] The effect is predominantly within (+0.241), with a reinforcing composition effect (+0.021).

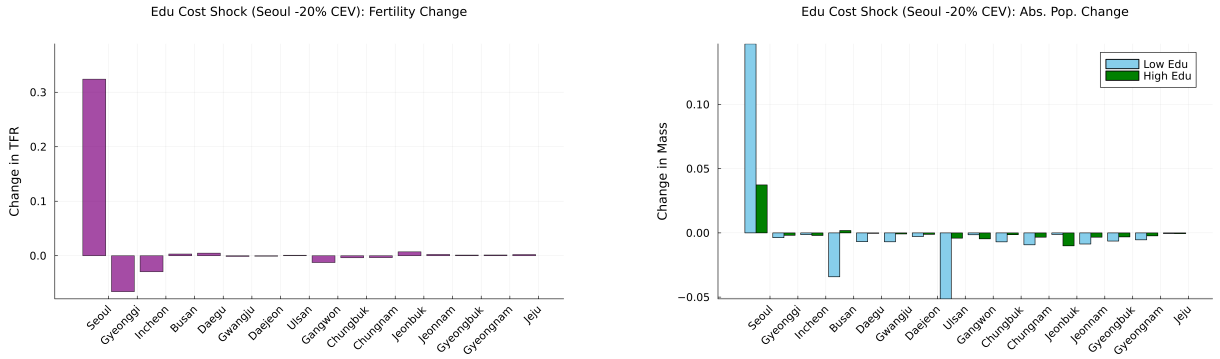
The positive within effect arises because nationwide education reform disproportionately benefits locations where education costs are highest, particularly Seoul and other metropolitan areas. By reducing the cost advantage of living outside the metro area, the reform narrows the fertility gap between Seoul and other regions. At the same time, some households on the margin between metropolitan and non-metropolitan locations shift toward regions where baseline fertility is higher, reinforcing the within effect rather than offsetting it as in the Seoul-only experiments.

The dominance of education cost reduction reflects the model's core mechanism. Education costs enter the household's problem through the cost function $\Phi_{j'}^{e,e',n'}$, which governs the price of investing in a child's education. This cost depends on the parent's education e, the target education level for the child e', the number of children n', and the neighborhood j'. Crucially, because Φ is increasing in n', families with more children face higher per-child education costs, formalizing the quantity-quality tradeoff. Reducing Φ lowers the marginal cost of maintaining education quality

[12] As noted in Section 7.1, percentage changes are reported as log differences. For effects of this magnitude, log and simple percentage changes diverge slightly.

as parity increases, making additional children less expensive in quality-adjusted terms. The effect is amplified in neighborhoods where high-skilled parents concentrate, because the local education benchmark is steeper there and the cost savings from reform are correspondingly larger.

Figure 7: Impact of Education Cost Reductions on Demographic and Spatial Distribution



(a) Regional Fertility Change

(b) Regional Population Change

Notes: The figure illustrates the spatial impacts of a 20% reduction in education costs (CEV) within the metropolitan area (Seoul). The left panel shows the resulting fertility increase, while the right panel highlights the migration-driven population growth.

Figure 7 shows that low-skilled workers are more sensitive to the education reform, exhibiting a stronger propensity to migrate toward Seoul compared to low-skilled workers. Nonetheless, the within effect is so large that it overwhelms the composition offset. This contrasts sharply with the commuting, housing, and amenity experiments, where the within effects are modest enough for composition to dominate.

7.4 Fiscal transfers and financing

We next consider subsidies to parents financed through two alternative tax instruments, a 20% tax on the income of nonparents and a revenue-neutral universal tax. Both raise fertility, but the choice of financing instrument matters substantially.

Tax on nonparents. The childless-tax-financed subsidy raises national fertility by 8.79% (Table 7). The effect is almost entirely behavioral (within = +0.129), with minimal spatial reallocation (composition = +0.008). The mechanism operates through the fertility choice in the model. Households choose between $n' = 0$ and $n' \geq 1$ by comparing utility without children ($U^{e,0}_{j',k'}$) against utility

with children plus the expected value of offspring net of fixed costs ($U^{e,n'}_{j',k'} + W^c - F^{e,n'}_{j'}$). A subsidy to parents raises the payoff for $n' \geq 1$, with the total subsidy increasing proportionally in n', while a tax on non-parents lowers the payoff for the $n' = 0$ state. Both mechanisms widen the utility gap between parenthood and childlessness, effectively shifting the optimal fertility choice toward having children.

Universal tax. The revenue-neutral universal tax generates a much smaller response (Table 7).[13] National fertility rises by only 2.83% (within = +0.040). The difference arises because the universal tax reduces disposable income for parents and nonparents alike, partially offsetting the subsidy's incentive effect. Under the childless tax, the full incidence falls on nonparents, which sharpens the relative payoff to having children.

The comparison between financing instruments has a clear policy implication. The total amount transferred to parents is the same in both cases, but the fertility response differs by a factor of three. This is because fertility responds to relative incentives, not just absolute transfer levels. Financing through a childless tax creates a stronger wedge between the utility of parenthood and childlessness than financing through a universal tax.

7.5 Coordinated policy packages

Single-margin interventions leave substantial gains unrealized. Education cost reduction addresses the quantity-quality wedge but does not change the relative attractiveness of parenthood at the extensive margin. Subsidies change the extensive margin but do not reduce the per-child cost of education competition. Combining the two targets both margins simultaneously, and because the policies operate on largely distinct channels, their effects are approximately additive.

Education reform with childless-tax-financed subsidy. This package raises national fertility by 21.69%, the largest effect in the model (Table 7). The decomposition shows a within effect of +0.336, a composition effect of +0.027. The total change (+0.362) is close to the sum of the education reform alone (+0.262) and the childless tax alone (+0.137), suggesting that the two

[13] We solve for the implied universal tax rate and the corresponding subsidy per child that ensures total tax revenue is equivalent to that generated under the childless tax policy. Our results indicate that the implied revenue-neutral universal tax rate is approximately 2.7 percent. This rate is significantly lower than the 20 percent childless tax rate, primarily due to the substantially broader tax base utilized in the universal scheme.

policies operate on largely distinct margins. The combined effect is slightly smaller than the sum of parts (0.362 vs. 0.399), indicating modest general equilibrium crowd-out, but the interaction is small relative to the total effect.

A notable feature of this experiment is that the composition effect is positive and sizable. This contrasts with the Seoul-only experiments, where composition was consistently negative. Recall that the composition term $\sum_j TFR_{j,0} \cdot \Delta s_j$ is positive when population shifts toward regions with higher baseline fertility. Under the combined policy, the nationwide education reform reduces the relative attractiveness of Seoul (where education costs were most burdensome), while the subsidy raises the payoff to parenthood everywhere. Together, these forces appear to redistribute population in a way that no longer concentrates it in the lowest-fertility region.

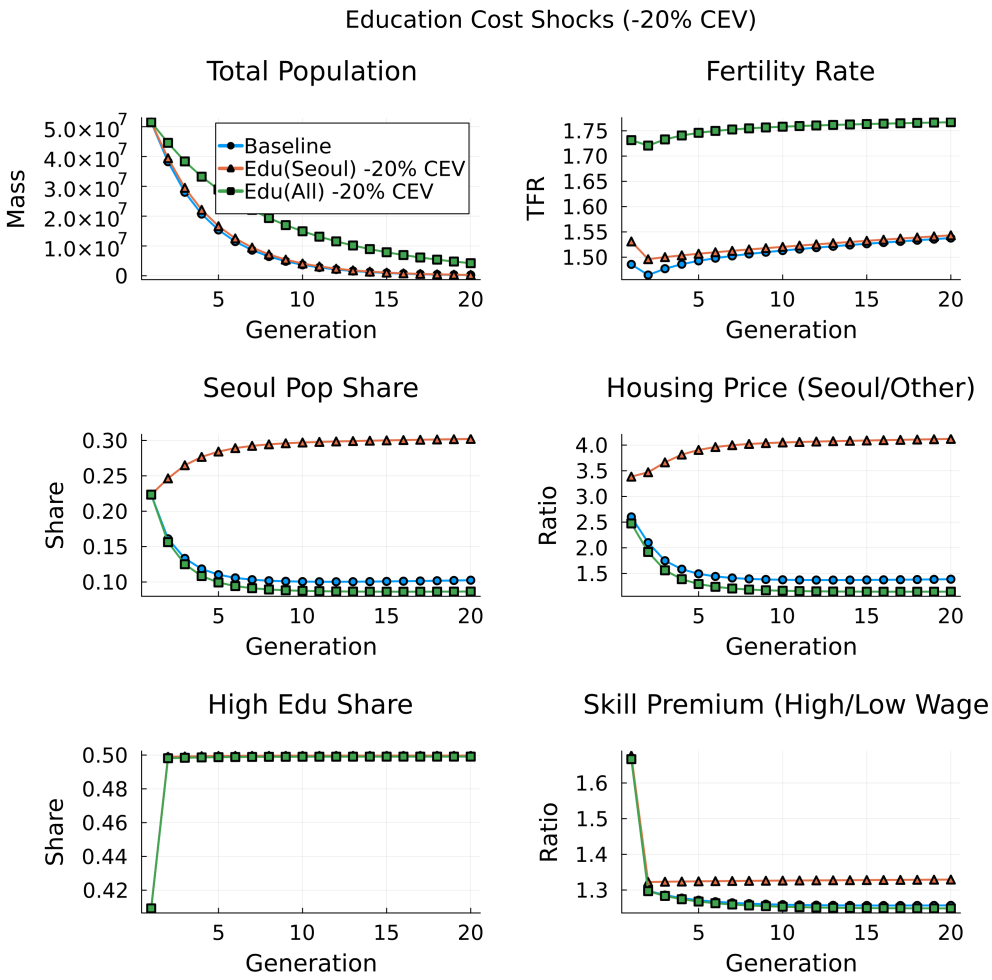
Education reform with universal-tax-financed subsidy. The same package financed by a universal tax generates a smaller but substantial response of 18.08% (Table 7). The within effect (+0.272) is lower than under the childless tax (+0.362), reflecting the weaker incentive from universal financing. The composition and covariance terms are similar.

7.6 Transition Dynamics

The counterfactual results reported above compare stationary equilibria in which young adults optimize given the existing old population. We complement these comparisons with transition paths that trace the economy's adjustment from the observed population distribution forward, allowing both the young and old populations to evolve endogenously across generations. Figures 8 and 9 show transition paths for the education cost and subsidy experiments, respectively.

Figure 8 displays the transition under education cost reduction. The fertility response is immediate and large. Under the nationwide reform, fertility jumps from roughly 1.50 to 1.67 in the first generation and continues rising toward approximately 1.70 as the population distribution adjusts. The Seoul-only reform produces a smaller initial increase to about 1.55. Seoul's population share reveals a sharp spatial adjustment in the early generations. Under the Seoul-only reform, Seoul retains a higher long-run population share (around 0.50) than under the baseline, reflecting the increased attractiveness of the metro area when education costs fall there. Under the nationwide reform, Seoul's share falls to levels similar to the baseline (around 0.24), because the reform removes

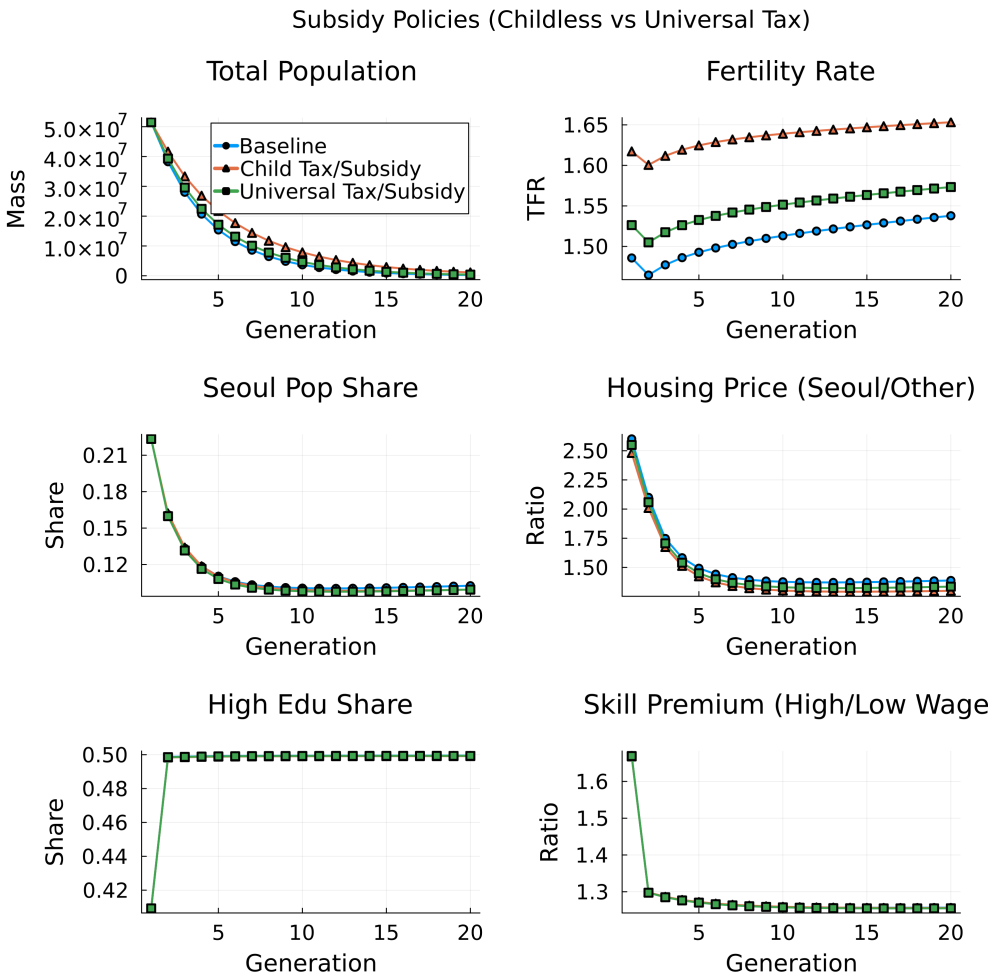
Figure 8: Transition with Education Reform



Notes: Each panel plots the transition path over 20 generations, starting from the observed population distribution. Blue circles denote the baseline (no policy), brown triangles denote the Seoul-only education cost reduction (-20% CEV), and green squares denote the nationwide education cost reduction (-20% CEV).

the relative cost disadvantage of living outside the capital. Relative housing prices follow a similar pattern. The Seoul-only reform sustains a long-run Seoul price premium roughly three times that of other regions, while the nationwide reform compresses relative prices toward parity. In all cases, most of the spatial reallocation occurs within the first three to four generations, with slower adjustment thereafter. Regardless of the education reform, the long-run fraction of college-educated children is anchored at approximately 45 percent, reflecting the high educational standards of Korean parents. In the long run, the skill premium stabilizes at around 2.7, a result driven by the significant oversupply of college-educated workers.

Figure 9: Transition with Child Subsidies



Notes: Each panel plots the transition path over 20 generations, starting from the observed population distribution. Blue circles denote the baseline (no policy), brown triangles denote the childless-tax-financed subsidy, and green squares denote the universal-tax-financed subsidy.

Figure 9 displays the transition under the subsidy experiments. The fertility effects are smaller in magnitude, consistent with the stationary comparisons in Table 7. The childless-tax-financed subsidy raises fertility from about 1.50 to roughly 1.60 in the first generation, converging toward approximately 1.63 by generation 20. The universal-tax subsidy produces an intermediate response. Unlike the education experiments, the subsidy policies have minimal effects on Seoul's population share or relative housing prices, confirming that their mechanism operates almost entirely through behavioral fertility responses rather than spatial reallocation.

7.7 Robustness Check and Additional Exercises

Figure 10 reports sensitivity of the main results along two dimensions. It traces the fertility response as a function of shock size for each policy instrument. The education cost reduction effect is approximately linear over the range considered, while housing supply, commuting cost, and amenity shocks produce smaller and flatter responses. This near-linearity suggests that the marginal fertility return to education reform does not diminish rapidly at larger shock sizes, reinforcing the finding that education costs are the binding margin.

Table A.6 presents our counterfactual exercises with alternative dispersion parameters: the dispersion parameter for idiosyncratic preference shocks for fertility choice (σ_N), the dispersion parameter for residential-commuting choice (σ_{JK}), and the dispersion of education preference shocks (σ_E). The results demonstrate that our main findings are quantitatively robust to these alternative parameterizations. While fertility tends to increase with larger values of these parameters, the relative ranking of policy effectiveness remains remarkably consistent across all specifications. In particular, the combination of education reform and subsidies consistently outperforms interventions such as housing supply or commuting cost reductions in raising the total fertility rate. This stability suggests that structural policies targeting education and deploying subsidies are fundamental to addressing fertility, regardless of any potential under- or over-estimation of the dispersion parameters.

Appendix H presents transition paths for each additional policy experiment. The primary takeaway is that various policy interventions prove effective in the long run as their effects accumulate across generations. As illustrated in Figure A.2, a 20% reduction in commuting costs (CEV) across all regions results in a significantly steeper increase in the total population after the tenth generation, reflecting the long-term compounding of the policy's impact on regional fertility and settlement patterns.

7.8 Discussion

Three findings from the counterfactual analysis deserve emphasis, as they bear on both the academic literature and the policy debate.

First, spatial general equilibrium can reverse the sign of a policy's effect. The Seoul-only

Figure 10: Policy Shock Analysis: Fertility Sensitivity and Tax Implications

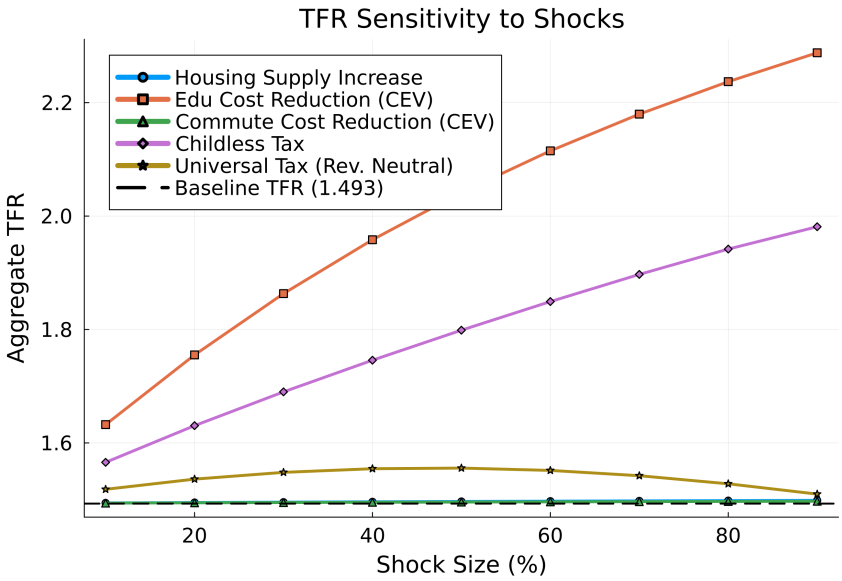


Figure 11: Sensitivity of Fertility to Shock Size

commuting, amenity, and TFP experiments all produce positive within-region fertility effects that are offset by negative composition effects. A partial equilibrium evaluation, or a reduced-form study that estimates the local treatment effect without accounting for migration responses, would conclude that these policies raise fertility. The model shows that they can lower it nationally. This result follows from the basic structure of a spatial equilibrium in which a low-fertility destination attracts population when it becomes more attractive, though its quantitative magnitude depends on the estimated parameters.

Second, education cost reduction is the most powerful single lever because it targets the structural driver of Korea's spatial fertility gradient. In locations where high-skilled parents concentrate, the local benchmark for education investment rises, and parents face a steeper quantity-quality tradeoff. Reducing education costs directly weakens this mechanism. No other single-margin intervention comes close, because housing, commuting, and amenity policies do not address the education competition externality.

Third, coordinated policies outperform isolated interventions. Combining education cost reduction with targeted subsidies generates a 13% fertility response, roughly the sum of the two individual effects. The policies target largely distinct margins, and their effects are approximately additive

with modest crowd-out, so pursuing both simultaneously captures gains that neither achieves alone.

These results connect to the broader literature on pronatalist policy. Reduced-form evaluations of baby bonuses and childcare subsidies typically find modest positive effects on fertility, consistent with our subsidy experiments (the revenue-neutral universal subsidy raises fertility by only 1.7%). Our contribution is to show that the modest effects are not evidence of policy ineffectiveness per se, but of targeting the wrong margin. When policy targets the structural source of low fertility, specifically the education competition externality that is amplified by spatial sorting, the response is several times larger.

8 Conclusion

This paper develops a dynamic spatial equilibrium model in which households jointly choose where to live, where to work, how much housing to consume, how many children to have, and how much to invest in their children's human capital, with wages and housing prices clearing in local markets. We estimate the model using administrative microdata from South Korea and use it to evaluate fertility policies through a spatial lens. The framework provides a unified account of why fertility is lowest in the highest-wage metropolitan areas and why large-scale pronatalist spending has failed to reverse the decline.

The central mechanism operates through spatial sorting and its downstream consequences. High-skilled workers concentrate in Seoul in pursuit of wage premia, but this concentration raises housing demand, lengthens commutes, and intensifies competition for educational status. The resulting increase in the cost of maintaining child quality sharpens the quantity-quality tradeoff, pushing fertility down precisely where incomes are highest. These forces interact in general equilibrium. A policy that makes Seoul more attractive to families also attracts more households into a low-fertility environment, and the compositional shift can offset or even reverse the intended behavioral effect.

The counterfactual experiments yield three findings with direct policy relevance. First, place-based interventions that target a single cost margin in the Seoul metropolitan area can backfire at the national level. Reducing commuting costs, expanding housing supply, improving amenities, or raising productivity in Seoul all produce positive within-region fertility effects, but the induced migration response reallocates population toward the metro area where baseline fertility is low.

The composition effect offsets the behavioral gain, leaving national fertility unchanged or slightly lower.

Second, nationwide education cost reduction is the most powerful single-policy lever, raising the national fertility rate by 23 percent when sized at 20 percent of consumption in welfare terms. This effect is several times larger than any other single-margin intervention because it directly weakens the education competition externality that spatial sorting amplifies. In locations where high-skilled parents concentrate, the local benchmark for educational investment rises and parents face a steeper quantity-quality tradeoff. Reducing education costs lowers the marginal cost of maintaining child quality as parity increases, and the effect is largest precisely in the locations where the externality is most severe.

Third, coordinated policies outperform any single intervention. Combining education reform with subsidies to parents financed by a tax on nonparents raises fertility by 27 percent, roughly the sum of the two individual effects. The two instruments target largely distinct margins, and their effects are approximately additive with modest crowd-out. This result suggests that Korea's sustained pronatalist spending has underperformed not because fiscal transfers are ineffective, but because transfers alone do not address the education competition externality. When transfers are combined with a policy that targets the structural source of low fertility, the response is substantially larger than either instrument achieves on its own.

These results carry a broader lesson for the design of pronatalist policy. Korea has spent over 280 trillion won on fertility programs since 2006, yet the total fertility rate has continued to fall. The model suggests that the limited effectiveness reflects not a lack of fiscal effort but a misalignment between the instruments deployed and the structural source of low fertility. Cash transfers and childcare subsidies lower the absolute cost of children, but they do not address the education competition externality that makes children progressively more expensive in locations where high-skilled parents concentrate. When policy targets the right margin, the response is substantially larger.

The analysis has several limitations that point toward future work. The model operates at the province level, which captures the key spatial margins but does not exploit within-province heterogeneity across neighborhoods. Extending the framework to finer geographic resolution would test whether the main quantitative results are robust to local variation in housing and education

markets.

The current specification treats education competition through a reduced-form cost function rather than modeling the strategic interaction among parents directly. A richer formulation with endogenous competition intensity could sharpen the welfare analysis and generate additional predictions about how interventions affect the equilibrium distribution of educational investment across locations.

The model also abstracts from gender-specific frictions in the labor market, including the career penalties associated with childbearing that feature prominently in the broader fertility literature. Incorporating these margins would allow the framework to speak to policies that reduce the opportunity cost of children through parental leave, flexible work arrangements, and workplace norms. In the Korean context, where female labor force participation remains well below the OECD average, these frictions likely interact with the spatial and education channels modeled here.

Finally, the transition dynamics indicate that the full effects of policy interventions accumulate over multiple generations as the population distribution adjusts. Understanding the speed of this adjustment and the role of expectations in shaping it would require a fully dynamic model with forward-looking agents, which we leave for future research.

Despite these limitations, the paper's core finding is clear. Fertility in South Korea is not simply low because children are expensive. It is low because the spatial concentration of high-skilled workers generates a set of reinforcing pressures through housing costs, commuting burdens, and education competition that standard policy instruments do not address. A framework that integrates these margins in general equilibrium is necessary both to understand why fertility has continued to fall and to identify the policy combinations most likely to reverse it.

References

Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics*, volume 4, pages 1043–1171. Elsevier.

Allen, T. and Arkolakis, C. (2014). Trade and the topography of the spatial economy. *Quarterly Journal of Economics*, 129(3):1085–1140.

Allen, T. and Arkolakis, C. (2025). Quantitative regional economics. In Donaldson, D. and Redding, S. J., editors, *Handbook of Regional and Urban Economics*, volume 6, pages 1–72. Elsevier.

Ang, G., Tan, Y., Zhai, Y., Zhang, F., and Zhang, Q. (2024). Housing wealth, fertility and children's health in China: A regression discontinuity design. *Journal of Health Economics*, 97:102915.

Arkolakis, C., Kwon, S., Russel, D., Shi, C., and Yang, H. (2025). Risk-seeking motives and speculative investment. Available at SSRN: `https://ssrn.com/abstract=5761702`.

Barro, R. J. and Becker, G. S. (1989). Fertility choice in a model of economic growth. *Econometrica*, 57(2):481–501.

Bartik, T. J. (1991). Who benefits from state and local economic development policies? Technical report, W.E. Upjohn Institute for Employment Research.

Bauernschuster, S., Hener, T., and Rainer, H. (2016). Children of a (policy) revolution: The introduction of universal child care and its effect on fertility. *Journal of the European Economic Association*, 14(4):975–1005.

Becker, G. S. (1960). An economic analysis of fertility. *Demographic and Economic Change in Developed Countries*, pages 209–240.

Becker, G. S. and Tomes, N. (1976). Child endowments and the quantity and quality of children. *Journal of Political Economy*, 84(4):S143–S162.

Bick, A. (2016). The quantitative role of child care for female labor force participation and fertility. *Journal of the European Economic Association*, 14(3):639–668.

Bokun, M. (2024). Does the family 500+ programme increase fertility? a cautionary tale for evidence-based policy. *Demographic Research*. Open-access journal article.

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

Choi, J. (2012). Private tutoring and educational inequality: Evidence from a dynamic model of academic achievement in Korea. `https://paa2013.populationassociation.org/papers/130384`. PAA Annual Meeting.

Coeurdacier, N., Combes, P.-P., Gobillon, L., and Oswald, F. (2023). Fertility, housing costs and city growth. Working paper.

Cohen, A., Dehejia, R., and Romanov, D. (2013). Financial incentives and fertility. *Review of Economics and Statistics*, 95(1):1–20.

Couillard, B. K. (2025). Build, baby, build: How housing shapes fertility. Working paper, University of Toronto, Department of Economics. Version: November 2, 2025.

Da Rocha, J. M. and Fuster, L. (2006). Why are fertility rates and female employment ratios positively correlated across o.e.c.d. countries? *International Economic Review*, 47(4):1187–1222.

Dahl, G. B., Løken, K. V., Mogstad, M., and Salvanes, K. G. (2016). What is the case for paid maternity leave? *Review of Economics and Statistics*, 98(4):655–670.

Dettling, L. J. and Kearney, M. S. (2014). House prices and birth rates: The impact of the real estate market on the decision to have a baby. *Journal of Public Economics*, 110:82–100.

Diamond, R. (2016). The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000. *American Economic Review*, 106(3):479–524.

Doepke, M., Hannusch, A., Kindermann, F., and Tertilt, M. (2023). The economics of fertility: A new era. In Lundberg, S. and Voena, A., editors, *Handbook of the Economics of the Family*, volume 1, pages 151–254. North-Holland, Amsterdam.

Doepke, M. and Kindermann, F. (2019). Bargaining over babies: Theory, evidence, and policy implications. *American Economic Review*, 109(9):3264–3306.

Eckert, F. and Kleineberg, T. (2024). The geography of opportunity: Education, work, and intergenerational mobility across US counties. Institute Working Paper, Federal Reserve Bank of Minneapolis.

Fazio, D., Ramadorai, T., Skrastins, J., and Van Doornik, B. F. N. (2024). Housing and fertility. SSRN Working Paper.

Fehr, H. and Ujhelyiova, D. (2013). Fertility, female labor supply, and family policy. *German Economic Review*, 14(2):138–165.

Gauthier, A. H. (2007). The impact of family policies on fertility in industrialized countries: A review of the literature. *Population Research and Policy Review*, 26(3):323–346.

Gauthier, A. H. and Hatzius, J. (1997). Family benefits and fertility: An econometric analysis. *Population Studies*, 51(3):295–306.

Goldin, C. (2021). *Career and Family: Women's Century-Long Journey toward Equity*. Princeton University Press, Princeton, NJ.

González, L. (2013). The effect of a universal child benefit on conceptions, abortions, and early maternal labor supply. *American Economic Journal: Applied Economics*, 5(3):160–188.

Guner, N., Kaya, E., and Sánchez-Marcos, V. (2024). Labor market institutions and fertility. *International Economic Review*, 65(3):1551–1587.

Havnes, T. and Mogstad, M. (2011). No child left behind: Subsidized child care and children's long-run outcomes. *American Economic Journal: Economic Policy*, 3(2):97–129.

Hoem, J. M. (1993). Public policy as the fuel of fertility: Effects of a policy reform on the pace of childbearing in sweden in the 1980s. *Acta Sociologica*, 36(1):19–31.

Hwang, J. (2023). Later, fewer, none? recent trends in cohort fertility in South Korea. *Demography*, 60(2):563–582.

Kearney, M. S. and Levine, P. B. (2025). Why is fertility so low in high income countries? Working Paper 33989, National Bureau of Economic Research.

Keng, S.-H., Lin, T.-C., and Sheu, H.-J. (2011). Can education subsidy reverse falling fertility? the case of taiwan. *The Journal of Development Studies*. Empirical evaluation of changes in child education subsidies.

Kim, S., Tertilt, M., and Yum, M. (2024). Status externalities in education and low birth rates in Korea. *American Economic Review*, 114(6):1576–1611.

Kim, W. (2024). Baby bonus, fertility, and missing women. *Journal of Human Resources*. Forthcoming.

Kleinman, B., Liu, E., and Redding, S. J. (2023). Dynamic spatial general equilibrium. *Econometrica*, 91(2):385–424.

Kravdal, Ø. (1996). How the availability of childcare affects fertility in norway. *Population Research and Policy Review*, 15(3):201–218.

Lalive, R. and Zweimüller, J. (2009). How does parental leave affect fertility and return to work? evidence from two natural experiments. *Quarterly Journal of Economics*, 124(3):1363–1402.

Laroque, G. and Salanié, B. (2014). Identifying the response of fertility to financial incentives. *Journal of Applied Econometrics*, 29(2):314–332.

Lovenheim, M. F. and Mumford, K. J. (2013). Do family wealth shocks affect fertility choices? evidence from the housing market. *Review of Economics and Statistics*, 95(2):464–475.

Milligan, K. (2005). Subsidizing the stork: New evidence on tax incentives and fertility. *Review of Economics and Statistics*, 87(3):539–555.

Ministry of Land, Infrastructure and Transport (2024a). 2024 housing survey. Technical report, Ministry of Land, Infrastructure and Transport, Sejong, Korea. Survey of 61,000 households. Median price-to-income ratio for Seoul reported as 13.9.

Ministry of Land, Infrastructure and Transport (2024b). Urban planning statistics []. Technical report, Ministry of Land, Infrastructure and Transport, Sejong, South Korea. National Approved Statistics No. 315002. Available at `https://www.eum.go.kr/web/cp/st/stUpisStatDet.jsp`.

OECD (2016). Time spent travelling to and from work. Technical Report LMF2.6, OECD Family Database. Korea: 58 minutes one-way (highest among OECD); OECD average: 28 minutes one-way.

OECD (2023). Pensions at a glance 2023: Country profile for korea. Technical report, OECD. Target replacement rate of 40% for 40 years of contributions, scheduled to take effect by 2028.

OECD (2024). Korea's unborn future. Technical report, Organisation for Economic Co-operation and Development. OECD Publishing, Paris.

Presidential Committee on Ageing Society and Population Policy (2022). The fourth basic plan for low fertility and aging society: 2022 implementation plan for central government ministries. Technical report, Government of the Republic of Korea, Seoul. In Korean.

Raute, A. (2019). Can financial incentives reduce the baby gap? evidence from a reform in maternity leave benefits. *Journal of Public Economics*, 169:203–222. First circulated as NBER Working Paper No. 23793 (2017).

Redding, S. J. and Rossi-Hansberg, E. (2017). Quantitative spatial economics. *Annual Review of Economics*, 9(1):21–58.

Rindfuss, R. R., Guilkey, D., Morgan, S. P., and Kravdal, Ø. (2010). Child-care availability and fertility in norway. *Population and Development Review*, 36(4):725–748.

Silva, J. M. C. S. and Tenreyro, S. (2006). The log of gravity. *The Review of Economics and Statistics*, 88(4):641–658.

Simon, C. J. and Tamura, R. (2009). Do higher rents discourage fertility? evidence from U.S. cities, 1940–2000. *Regional Science and Urban Economics*, 39(1):33–42.

Slonimczyk, F. and Yurko, A. (2014). Assessing the impact of the maternity capital policy in russia using a dynamic model of fertility and employment. IZA Discussion Paper 8401, IZA Institute of Labor Economics.

Sommer, K. (2016). Fertility choice in a life cycle model with idiosyncratic uninsurable earnings risk. *Journal of Monetary Economics*, 83:27–38.

Statistics Korea (2016). Final results of birth statistics in 2015. Technical report, Statistics Korea.

Statistics Korea (2023). Commute time statistics. Statistics Korea, using SK Telecom mobile tracking data. National average: 72.6 minutes round-trip; Seoul Metropolitan Area: 83.2 minutes. Reported December 2023.

Statistics Korea (2024). Birth statistics in 2023. Technical report, Statistics Korea.

Stock, J. H. and Yogo, M. (2005). Testing for weak instruments in linear IV regression. In Andrews, D. W. K. and Stock, J. H., editors, *Identification and Inference for Econometric Models*, pages 80–108. Cambridge University Press.

Appendix

A Institutional Background

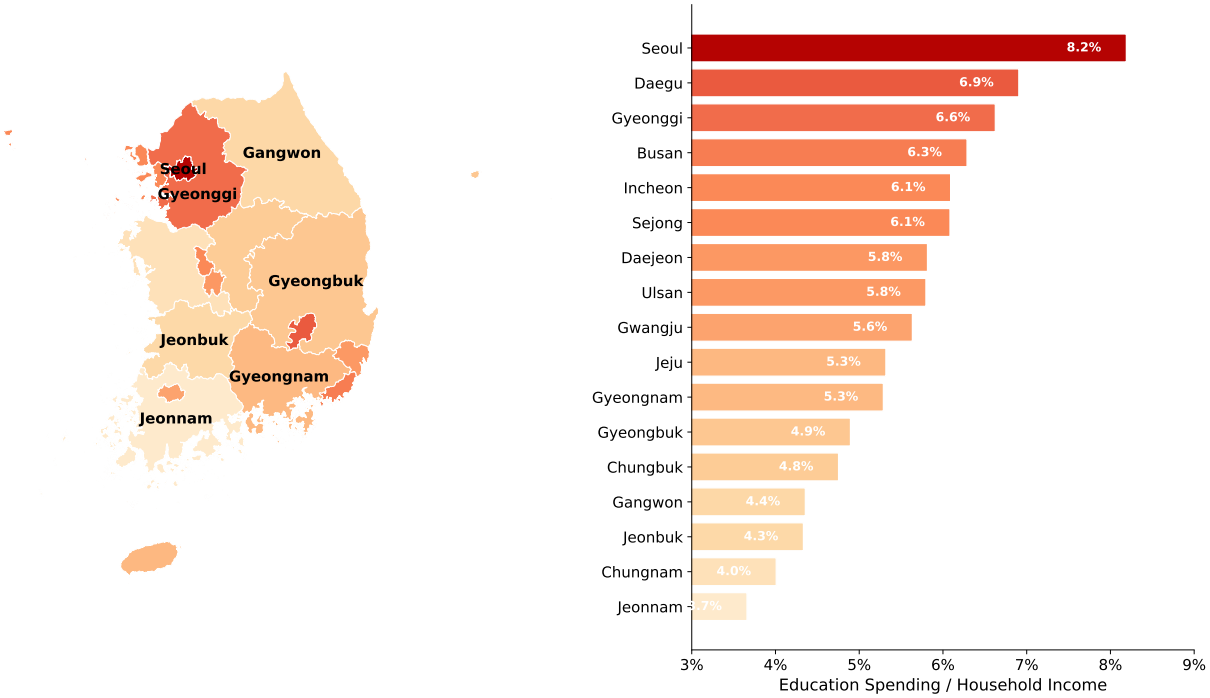
The model requires households to make joint decisions over residential location, workplace, housing, and investment in children. This section describes the Korean institutional environment that shapes these choices and the data infrastructure that makes estimation possible.

A.1 Education

Education is the highest-stakes investment Korean parents make, and the system that determines its returns is unusually narrow (Kim et al., 2024). Nearly all university placement depends on a single national exam, the *suneung*, administered once per year. Performance on this exam determines access to elite universities, which shapes lifetime earnings, marriage market outcomes, and social standing. The consequences of falling short are severe and widely understood, so parents begin investing early and continue through high school.

This pressure fuels the world's largest private tutoring industry (Choi, 2012). Private academies known as *hagwon* offer after-school instruction in every subject, and household spending on *hagwon* constitutes a major budget item. But access to high-quality tutoring is not uniform. Seoul concentrates the most competitive and expensive academies, while rural provinces have far thinner markets and lower per-child expenditures. Figure A.1 illustrates this disparity. Households in Seoul spend 8.2 percent of their income on private education per child, more than twice the 3.7 percent observed in Jeonnam. The model captures this variation through location-specific education costs that depend on the local skill composition, generating a quantity-quality tradeoff that becomes steeper where high-skilled parents are concentrated. Direct spending on private academies, however, understates the full cost of education competition. Time spent driving children between multiple hagwon sessions, opportunity costs of parental involvement, and the competitive pressure to match local achievement norms all raise the effective cost of maintaining child quality beyond what spending data capture.

Figure A.1: Education spending intensity by province



Notes: Median ratio of annual private education spending per child to annual household income, by province, pooled over 2020–2024. Data are from the Private Education Expenditures Survey, jointly administered by Statistics Korea and the Ministry of Education. Seoul exhibits the highest spending intensity at 8.2 percent, while Jeonnam records the lowest at 3.7 percent.

A.2 Housing

Housing prices vary sharply across Korean regions, with Seoul several times more expensive than the national average. The median Seoul home costs roughly 14 years of household income, making it one of the least affordable markets in the developed world (Ministry of Land, Infrastructure and Transport, 2024a). This concentration of high prices reflects both the location of high-wage employment in the capital region and regulatory constraints that limit new construction. Seoul's greenbelt, established in 1971, restricts development on roughly a quarter of the city's land area (Ministry of Land, Infrastructure and Transport, 2024b), and reconstruction of aging apartment complexes requires supermajority approval from existing owners, a process that often takes a decade or longer. In the model, housing supply is fixed at observed levels, so increased demand from sorting translates directly into higher prices.

To afford housing in expensive areas, households take on substantial debt. The ratio of household debt to disposable income exceeds 200 percent, among the highest in the OECD (Statistics

Korea, 2024). Much of this borrowing finances *jeonse*, a rental system in which tenants deposit a lump sum of 50 to 80 percent of the property's value instead of monthly rent. For young families considering children, these costs directly compete with resources available for child-rearing. In the model, where housing prices clear in local markets, the same sorting forces that draw workers to Seoul also put upward pressure on its housing costs.

A.3 Commuting

Korean workers commute longer than almost anywhere else in the developed world (OECD, 2016). The average worker's commute is 73 minutes per day, and in the Seoul Metropolitan Area it reaches 83 minutes (Statistics Korea, 2023). These long commutes reflect a labor market in which high-wage jobs cluster geographically and workers live farther out where housing is more affordable. Commuting thus allows households to access Seoul wages without paying Seoul housing prices. But time spent commuting is time unavailable for childcare, and long commutes may affect fertility decisions directly. The model treats commuting and migration as distinct choices. Households trade off longer commutes against higher housing costs, and both margins affect the resources available for children.

A.4 Policy Environment

Over the past two decades, Korean policymakers have tried to raise fertility by offering baby bonuses, childcare subsidies, parental leave expansions, and preferential housing programs for families with children (OECD, 2024). Cumulative spending on these programs exceeded 280 trillion won ($200 billion) between 2006 and 2021 (Presidential Committee on Ageing Society and Population Policy, 2022). Yet fertility has continued to decline, falling from 1.24 in 2015 to 0.72 in 2023 (Statistics Korea, 2016, 2024).

This persistent drop in fertility despite large-scale intervention is the central policy puzzle. If the problem were simply that children are expensive, cash transfers should help. If the problem were inadequate childcare, subsidies should help. The fact that fertility keeps falling suggests that no single margin is the binding constraint. Instead, policies targeting one margin may be partially offset by equilibrium adjustments in others. For example, a childcare subsidy that makes Seoul more attractive to young families also bids up housing prices and intensifies education competition.

The model is designed to capture these interactions, which we explain further in counterfactual policy evaluations in Section 7.[14] In evaluating these programs, a key question is how the financing instrument shapes incentives. A subsidy to parents financed by a broad-based tax changes the absolute cost of children, while one financed by a tax falling specifically on nonparents also changes the relative payoff of parenthood versus childlessness. Section 7 uses both financing structures as theoretical benchmarks to isolate this distinction.

B Details on data processing

B.1 House Price

We obtain apartment transaction data from the Ministry of Land, Infrastructure, and Transport. Apartment prices are divided by floor area and adjusted for CPI to obtain prices per square meter. District-level prices are computed as the average transaction price in 2020.

B.2 Population & Birth

We obtain population data by age and location and birth data from Ministry of the Interior and Safety. The working-age population is defined as individuals aged 15–64, minors as those aged 0–14, and the elderly population as those aged 65 and above.

B.3 Labor Supply and Wage

We obtain data from the Local Area Labor Force Survey provided by Statistics Korea (MDIS). We pool the first- and second-half surveys. A worker is classified as high-skilled if they have graduated from a four-year university. Labor supply is measured as the total number of workers in the survey, and wages are defined as the average wage of workers aged 25–54.

B.4 Consumption Share on Rent

We obtain data from the Housing Survey conducted by Statistics Korea (MDIS). We compute imputed rent for deposit-based housing and owner-occupied housing using a 4.5% rate. Housing

[14]The national pension system also enters the household's intertemporal budget constraint, targeting a replacement rate of roughly 40 percent of career average earnings for workers with full contribution histories (OECD, 2023). We specify the pension system in Section 4.

expenditure shares by education level and number of children are defined as the median values across households. Due to the small number of observations, we impute values for households with four or more children using the value for four-child households.

B.5 Geographic Data

We obtain geographic shape file data from Ministry of the Interior and Safety. The distance between two cities (defined at the 4-digit level) is calculated as the straight-line distance between their centroids. The distance between two SiDo regions is defined as the average of the pairwise distances between all 4-digit units belonging to each region. School location data are obtained from the Ministry of Education. Address information is converted into geographic coordinates using Kakao Map, a widely used digital map in Korea.

B.6 Education

We obtain data on academic achievement from ASIL, which is based on raw data provided by the Ministry of Education, and use it in descriptive linear regressions. We use the Private Education Expenditure Survey from Statistics Korea to estimate education costs in the structural model. Using the Population and Housing Census, we also obtain information on the highest educational attainment of household members. Based on these Census data, Table A.1 reports the joint distribution of education levels within couples.

Table A.1: Educational Composition of Couples

	Share
Both college educated	0.33
Both non-college educated	0.43
Different education levels	0.25

Notes: Data are from the 2 percent sample of the 2020 Population and Housing Census. The sample is restricted to married couples both aged 25–55. Couples are classified into three groups: both college educated, both non-college educated, and mixed-education couples, based on completion of a four-year college degree.

B.7 Move and Birth Share

We calculate household-level moving and birth shares based on the mother. A woman is defined as the mother if she is identified as the household head or the spouse of the household head in the Census. We classify her as high educated if she has completed a four-year college degree and as low educated otherwise

B.8 Daycare

We obtain daycare data from the Ministry of Education and count the number of daycare centers whose observation date falls between their opening and closing dates.

B.9 Location Code Crosswalk

Location codes from Statistics Korea differ from administrative neighborhood codes. We construct a crosswalk by matching SiDo and SiGu names.

C Estimation: Education Costs

C.1 Analytical Framework

Recall that the probability $\pi_{j'}^{e,e',n'}$ that a parent with education level e in neighborhood j' chooses education level e' for their n' children is given by:

$$\pi_{j'}^{e,e',n'} = \frac{\exp\left(-\phi_{j'}^{e,e',n'} + \beta^y v^y(e', j')\right)}{\sum_{e''} \exp\left(-\phi_{j'}^{e,e'',n'} + \beta^y v^y(e'', j')\right)} \tag{A.1}$$

Here, by defining the normalized cost as $\phi = \Phi/\sigma_E$ and the normalized value as $v^y = V^y/\sigma_E$, we simplify the expression (10) for estimation. The denominator of this choice probability reflects a parent-neighborhood-specific inclusive value, conditional on the number of children. This value determines the expected return on children's education, $w^c (\equiv W^c/\sigma_E)$, which is subsequently utilized in the household's fertility decision:

$$w^c = \log\left(\sum_{e''} \exp(-\phi_{j'}^{e,e'',n'} + \beta^y v^y(e'', j'))\right) \tag{A.2}$$

D Estimation: Commuting Costs

D.1 Derivation of commuting odd difference

Rewriting the value of young adults when choosing the number of children in equation (3), we have:

$$V^{y,c}(e, j', k') = -C_{j,k'} + \log((1-\tau_p)Y_{k'}^e) + \mathbb{E}\left[\max_{n'}\{\sigma_N \varepsilon'_n + \mathbb{I}(n' \geq 1)(W^c(e, j', n') - F_{j'}^{e,n'}) + A_j^{e,n'} - \alpha^{e,n'} \log p_j\}\right] \tag{A.3}$$

Taking the difference between the value at commuting zone k' and k'', we obtain:

$$V^{y,c}(e, j', k'') - V^{y,c}(e, j', k') = \log Y_{k''}^e - \log Y_{k'}^e - (C_{j,k''} - C_{j,k'}) \tag{A.4}$$

From equation (26), this leads to the following log-relative share:

$$\log\left(\frac{\lambda^e_{j,j',k''}}{\lambda^e_{j,j',k'}}\right) = \frac{1}{\sigma_{JK}}\left[\log\left(\frac{Y^e_{k''}}{Y^e_{k'}}\right) - (C_{j,k''} - C_{j,k'})\right] \tag{A.5}$$

D.2 Instrumental Variable Strategy

Time differencing. Taking first differences of the log commuting odds ratio eliminates time-invariant location amenities and commuting cost components:

$$\Delta \log \frac{\lambda^e_{j,j',k'',t}}{\lambda^e_{j,j',k',t}} = \frac{1}{\sigma_{JK}} \Delta(g^e_{k'',t} - g^e_{k',t}) \tag{A.6}$$

where $g^e_{k,t} \equiv \log Y^e_{k,t} - \log Y^e_{k,t-1}$ denotes log wage growth in commuting zone k for education group e between periods $t-1$ and t.

Bartik instrument. The instrument for wage growth in zone k is constructed as

$$B^e_{k,t} = \sum_i s^e_{i,k,2010} \cdot \Delta \log \bar{w}^e_{-k,i,t} \tag{A.7}$$

where $s^e_{i,k,2010}$ is the employment share of industry i among workers with education e in zone k in 2010, and $\Delta \log \bar{w}^e_{-k,i,t}$ is the leave-one-out national log wage growth in industry i for education group e, computed excluding zone k to avoid mechanical correlation. The sum runs over 77 two-digit industries classified under the Korean Standard Industrial Classification (KSIC, 10th revision). The instrument for the pairwise wage growth differential is $B^e_{k'',t} - B^e_{k',t}$.

Identifying assumption. Under the framework of Borusyak et al. (2022), identification requires that national industry-level wage growth is uncorrelated with local commuting demand shocks, conditional on the initial industry composition. This is plausible because national industry wage growth in Korea is largely driven by export demand and global commodity cycles rather than local labor market conditions. The leave-one-out construction ensures that zone k's own wages do not enter its instrument.

Sample construction. The estimation sample pools two time periods: 2010–2015 and 2015–2020. Commuting zones are defined at the province (*sido*) level, yielding 15 zones. For each (education,

year, origin) cell, we form all pairwise combinations of destination zones, imposing a canonical ordering ($k' < k''$) to eliminate mirrored pairs and dropping self-pairs ($k' = k''$). Observations with missing commuting flows are excluded. The final sample contains 1,749 observations.

Estimation and first-stage diagnostics. Estimation is by two-stage least squares (2SLS), with standard errors clustered at the origin province level (17 clusters). The first-stage regression of the pairwise wage growth differential on the Bartik instrument differential yields a coefficient of 0.799 (robust SE = 0.179). The Kleibergen-Paap F-statistic under origin clustering is 32.3, exceeding the Stock and Yogo (2005) critical value of 16.38 for 10 percent maximal IV size. Under heteroskedasticity-robust standard errors without clustering, the Kleibergen-Paap F-statistic is 19.8. The Anderson-Rubin 95 percent confidence interval under origin clustering is $[0.4, 15.6]$, confirming that the positive effect is robust to weak-instrument concerns.

D.3 Estimation of Commuting Cost

Given the IV estimate $1/\hat{\sigma}_{JK} = 8.14$ (implying $\hat{\sigma}_{JK} = 0.12$), we recover commuting costs by assuming a linear relationship with distance: $C_{j',k} = \xi \cdot \text{dist}_{j',k}$. Substituting the estimated $\hat{\sigma}_{JK}$ into equation (19) and netting out the wage component, we regress the residualized log commuting ratio on the distance differential by OLS to obtain the estimate $\hat{\xi}$. Then, we have the estimates for commuting costs using $\hat{C}_{j',k} = \hat{\xi} \cdot \text{dist}_{j',k}$.

Table A.2: Regression for Commuting Cost

Dependent Variable:	$\log \frac{\lambda^e_{j,j',k''}}{\lambda^e_{j,j',k'}} / \frac{1}{\sigma_{JK}} - \log \frac{Y^e_{k''}}{Y^e_{k'}}$
$dist_{j'k''} - dist_{j'k'}$	-0.0047***
	(0.0003)
Observations.	4,548

Notes: OLS estimates. The dependent variable is the log commuting odds ratio scaled by the IV estimate $1/\hat{\sigma}_{JK}$ from Table 5, net of the log wage ratio. The regressor is the difference in distance (km) between the residence and each commuting zone in the pair. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

E Estimation: Local Amenities

E.1 Estimation of birth dispersion parameter σ_N

For households with no children ($n' = 0$), the birth rate is determined by the ratio of local amenities to commuting costs:

$$\gamma^{e,0}_{j',k'} = \frac{\exp(u^{e0}_{j'k'})}{\exp(v^{y,c}(e,j',k')\sigma_{JK}/\sigma_N)} \tag{A.8}$$

The decision rule for non parents is rewritten as:

$$\begin{aligned}\log \gamma^{e,0}_{j',k'} &= u^{e,n'}_{j'k'} - \frac{\sigma_{JK}}{\sigma_N} v^{y,c}(e,j',k') \\ &= a^{e,n'}_{j'} + \frac{1}{\sigma_N}[\log(\frac{Y^e_k}{p^{\alpha^{e,0}}_{j'}}) - \sigma_{JK} v^{y,c}(e,j',k') - C_{j'k'}]\end{aligned}$$

For households with children ($n' > 0$), the birth rate is given by:

$$\gamma^{e,n'}_{j',k'} = \frac{\exp(u^{e,n'}_{j'k'} + w^c(e,j',n')\sigma_E/\sigma_N - f^{e,n'}_{j'})}{\exp(v^{y,c}(e,j',k')\sigma_{JK}/\sigma_N)} \tag{A.9}$$

Convert the birth decision rule above into the following log formula:

$$\begin{aligned}\log \gamma^{e,n'}_{j',k'} &= u^{e,n'}_{j'k'} + \frac{\sigma_E}{\sigma_N} w^c(e,j',n') - \frac{\sigma_{JK}}{\sigma_N} v^{y,c}(e,j',k') - f^{e,n'}_{j'} \\ &= \underbrace{[a^{e,n'}_{j'} - f^{e,n'}_{j'} + \frac{\sigma_E}{\sigma_N} w^c(e,j',n')]}_{\text{Edu}\times\text{Neighborhood}\times\text{Number of children FE}} + \frac{1}{\sigma_N}[\log(\frac{Y^e_k}{p^{\alpha^{e,n'}}_{j'}}) - \sigma_{JK} v^{y,c}(e,j',k') - C_{j'k'}]\end{aligned}$$

Commuting costs $c_{j',k'}$, taste shock parameter σ_{JK}, and $v^{y,c}(e,j',k')$ have already been obtained from previous estimation step in Appendix D. Using regression below, we estimate $\hat{\sigma}_N (= 1/1.339 \approx 0.75)$ and the fixed effect term, $a^{e,n'}_{j'} - f^{e,n'}_{j'} + \frac{\sigma_E}{\sigma_N} w^c(e,j',n')$.

E.2 Estimation of Amenities

We identify the net amenity costs, $a^{e,n'}_{j'} - f^{e,n'}_{j'}$, through the following procedure. First, the value of child education, w^c, is estimated as a province-level fixed effect–interacted with parental education–derived from the education cost regression in (16). Given the estimated birth taste shock parameter σ_N, we implement a nested calibration for the education taste shock parameter σ_E. For each

Table A.3: Amenity Regression

Dependent Variable:	Birth Rate ($\gamma^{e,n'}_{j',k'}$)
$\log(Y/(P^{\alpha})) - \sigma_{JK} v^{y,c} - C_{j'k'}$	1.339**
	(0.497)
Observations.	527

Notes: Education×neighborhood×children fixed effects, income-to-housing prices, and neighborhood×commuting zone fixed effects (commuting costs) are controlled. Standard errors in parentheses are clustered at the province level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

candidate value of σ_E, we extract the net amenity term from the estimated fixed effect, defined as $FE \equiv a^{e,n'}_{j'} - f^{e,n'}_{j'} + \frac{\sigma_E}{\sigma_N} w^c(e, j', n')$. Using these recovered terms, we solve for the competitive stationary equilibrium of the model. We then calibrate σ_E to match the unconditional variance of education outcomes ($\pi^{e,e',n'}_{j'}$) observed in the Private Education Survey via a method of simulated moments. This iterative process yields the final calibrated σ_E and the corresponding structural net amenity costs.

F Model Solution Algorithm

F.1 Overview

The model is solved using a nested fixed-point algorithm that computes the stationary general equilibrium. The algorithm iterates through three layers:

1. **Inner Loop:** Solves for the value functions (V^y, V^{yc}, V^o, W^c) given a guess of prices and wages.

2. **Middle Loop:** Computes the stationary population distribution given the optimal decision rules derived from the value functions.

3. **Outer Loop:** Updates housing prices and wages to clear markets, iterating until convergence.

Algorithm 1 Value Function Iteration (VFI) for Stationary Equilibrium

Initialize parameters Θ, guess prices $p^{(0)}$ and wages $Y^{(0)}$.
Set iteration counter $t = 0$.
while convergence criterion $> \epsilon_{GE}$ **do**
    $t \leftarrow t + 1$
— **Step 1: Inner Loop (Value Functions)** —
    Initialize $V^{y,(0)}$.
    Compute Old Age Value $V^o(e, j)$ (static given $p^{(t-1)}$).
    Guess initial $V^{y,(0)}$.
    Compute Old Age Value $V^o(e, j)$ (static given prices $p^{(t-1)}$ from previous outer loop iteration).
    repeat
        Compute Child Quality Value $W^c(e, j, n)$ using backward induction on education choice.
        Compute Fertility Choice Value $V^{yc}(e, j, k)$ (Log-sum-exp over n).
        Compute Location Choice Value $V^{y,new}(e, j)$ (Log-sum-exp over j', k').
        if Endogenous Subsidy **then**
            Update subsidy rate τ_{sub} to balance budget given current choice probabilities.
            Update subsidy amount S_{kid} to balance budget given current choice probabilities.
        end if
        Update $V^y \leftarrow V^{y,new}$.
    until $||V^y - V^{y,old}|| < \epsilon_V$
— **Step 2: Middle Loop (Stationary Distribution)** —
    Construct transition matrix $\Lambda(e, j, j', k)$ from policy functions.
    Initialize population distribution $\Pi^{(0)}(e, j)$.
    Guess initial population distribution $\Pi^{(0)}(e, j)$.
    repeat
        $\Pi^{(k+1)}(e, j') = \sum_j \Pi^{(k)}(e, j) \sum_k \Lambda(e, j, j', k)$
    until $||\Pi^{(m+1)} - \Pi^{(m)}|| < \epsilon_{dist}$
    $\Pi^{(m+1)}(e, j') = \sum_j \Pi^{(m)}(e, j) \sum_k \Lambda(e, j, j', k)$
    Compute Mass $L^y(e, j, k) = \sum_{j_{prev}} \Pi(e, j_{prev}) \times \Lambda(e, j_{prev}, j, k)$ (stationary mass of young).
— **Step 3: Outer Loop (Market Clearing)** —
    Update Wages:
    Calculate aggregate labor supply $L^e_k = \sum_j L^y(e, j, k)$.
    Compute new wages $Y^{new}(e, k)$ from CES production function FOCs.
    Update Housing Prices:
    Calculate total housing expenditure $E_j = \sum_{e,k} L(e, j, k) \times \text{HousingShare} \times \text{Income}$.
    Compute new prices $p^{new}_j = E_j / H^{supply}_j$.
— **Convergence Check** —
    Compute distance $dist = \max(||p^{new} - p^{(t-1)}||, ||Y^{new} - Y^{(t-1)}||)$.
    Update with dampening:
    $p^{(t)} = \lambda p^{(t-1)} + (1 - \lambda) p^{new}$
    $Y^{(t)} = \lambda Y^{(t-1)} + (1 - \lambda) Y^{new}$
    $Y^{(t)} = \lambda Y^{(t-1)} + (1 - \lambda) Y^{new}$ ▷ $(t-1)$ denotes value from previous iteration
end while
return Equilibrium Prices, Distributions, and Value Functions.

F.2 Key Equations

Value Functions. The value functions are solved via backward induction within the period:

$$V^o(e,j) = A_{old}(e,j) + \ln(b(e,j)) - \alpha_{old}(e)\ln(p_j)$$

$$\text{where } b(e,j) = \rho_p \bar{Y}_{young}(e,j) \text{ (Pension Income)}$$

$$W^c(e,j,n) = \sum_{i=1}^{n} \sigma_E \ln\left(\sum_{e'} \exp\left(\frac{-\Phi_{e,e'}(i,j) + \beta_y V^y(e',j)}{\sigma_E}\right)\right)$$

$$V^{yc}(e,j,k) = \sigma_N \ln\left(\sum_{n} \exp\left(\frac{U(e,j,k,n) + W^c(e,j,n)}{\sigma_N}\right)\right)$$

$$V^y(e,j) = \sigma_{JK} \ln\left(\sum_{j',k'} \exp\left(\frac{-T^e_{jj'} + \beta_c V^{yc}(e,j',k') + \beta_o V^o(e,j')}{\sigma_{JK}}\right)\right)$$

Market Clearing. Labor Market: Wages are determined by the marginal product of labor in a CES production function with elasticity of substitution ρ:

$$Y^e_k = Z_k S^e_k (L^e_k)^{-\frac{1}{\rho}} \left(\sum_{e'} S^{e'}_k (L^{e'}_k)^{\frac{\rho-1}{\rho}}\right)^{\frac{1}{\rho-1}}$$

Housing Market: Households have Cobb-Douglas preferences over consumption and housing with expenditure share $\alpha(e,n)$. This yields the indirect utility term $\ln(I(e,k,n)) - \alpha(e,n)\ln(p_j)$ and individual housing demand $h = \frac{\alpha(e,n)I(e,k,n)}{p_j}$.

$I(e,k,n)$ denotes net income after taxes and subsidies:

$$I(e,k,n) = \begin{cases} Y^e_k(1-\tau_p)(1-\tau_{univ})(1-\tau_{child}) & \text{if } n = 0 \\ Y^e_k(1-\tau_p)(1-\tau_{univ}) + n \cdot S_{kid} & \text{if } n > 0 \end{cases}$$

Prices clear when total demand equals fixed supply H_j:

$$p_j = \frac{1}{H_j}\left(\sum_{e,k} L^y(e,j,k) \sum_{n} P(n|e,j,k)\alpha(e,n)I(e,k,n) + \sum_{e} L^o(e,j)\alpha(e,0)Y^o(e,j)\right)$$

Numerical Stability: Log-Sum-Exp Trick. The value functions involve terms of the form $\sigma \ln\left(\sum \exp(V_i/\sigma)\right)$. Direct computation can lead to numerical overflow if V_i/σ is large. To ensure

stability, we use the Log-Sum-Exp identity:

$$\ln\left(\sum_i \exp(x_i)\right) = M + \ln\left(\sum_i \exp(x_i - M)\right), \quad \text{where } M = \max_i x_i$$

By shifting the exponents by the maximum value M, we ensure that the largest term in the sum is $\exp(0) = 1$ and all other terms are ≤ 1, preventing overflow.

F.3 Counterfactual Analysis

To assess the impact of policy changes or economic shocks, we perform counterfactual experiments. The procedure is as follows:

1. **Baseline Calibration:** Start with the model parameters Θ_{base} calibrated to match the observed data moments. Let $\mathbf{X}_{base}$ be the vector of equilibrium outcomes (e.g., TFR, population distribution, prices).

2. **Parameter Modification:** Introduce a shock by modifying a subset of parameters to Θ_{cf}. For example, where δ denotes the shock size:

   - *Housing Supply Shock:* $H_j^{cf} = (1+\delta)H_j^{base}$.
   - *Education Cost Shock:* $\Phi_{cf} = \Phi_{base} - \ln(1+\delta)$ (Consumption Equivalence Variation).
   - *Subsidy Policy:* We introduce a subsidy S_{kid} per child funded by a tax τ. The government budget constraint is solved endogenously:

     $$S_{kid} \sum_{e,j,k} L^y(e,j,k)\mathbb{E}[n|e,j,k] = \text{Total Revenue}$$

     where revenue depends on the tax base:

     - **Childless Tax:** τ_{child} applies to households with $n = 0$. Revenue $= \tau_{child} \sum_{e,j,k} L^y(e,j,k)P(n = 0|e,j,k)Y_k^e(1-\tau_p)$.
     - **Universal Tax:** τ_{univ} applies to all households. Revenue $= \tau_{univ} \sum_{e,j,k} L^y(e,j,k)Y_k^e(1-\tau_p)$.

3. **Re-solving Equilibrium:** Solve for the new stationary general equilibrium using the algorithm described in Section 2, yielding outcomes $\mathbf{X}_{cf}$.

4. **Comparison and Decomposition:** Compare $\mathbf{X}_{cf}$ against $\mathbf{X}_{base}$. The change in the aggregate Total Fertility Rate (TFR) is decomposed into three components:

$$\Delta \text{TFR} = \underbrace{\sum_j \Delta \text{TFR}_j \cdot s_{j,base}}_{\text{Within Effect}} + \underbrace{\sum_j \text{TFR}_{j,base} \cdot \Delta s_j}_{\text{Composition Effect}} + \underbrace{\sum_j \Delta \text{TFR}_j \cdot \Delta s_j}_{\text{Covariance}}$$

where s_j is the population share of region j, $\Delta \text{TFR}_j = \text{TFR}_{j,cf} - \text{TFR}_{j,base}$, and $\Delta s_j = s_{j,cf} - s_{j,base}$.

F.4 Transition Dynamics

To analyze the dynamic trajectory of the economy following a shock, we simulate the transition path from the initial state (calibrated to observed data) to the new counterfactual steady state. Notably, even in the baseline scenario without an external shock, we observe transitional dynamics as the economy converges toward its long-run steady state from the initial observed data point, which is not yet in a steady state.

Nature of the Shock and Expectations. The exercise is modeled as an *MIT Shock*, representing an unexpected and permanent change in parameters at $t = 1$. Agents immediately observe these new parameters and form expectations based on a "Sudden Expectations Adjustment": they assume the economy has instantaneously jumped to the new steady state for evaluating future continuation values, using the new steady-state value functions (V_{SS}, W^c_{SS}) for forward-looking decisions such as migration and education investment. Consequently, agents hold myopic price expectations, acting as if future prices are fixed at the new steady state levels, although their current decisions remain constrained by the actual market-clearing prices p_t and Y_t that evolve dynamically as the population distribution L_t converges.

Algorithm. Based on the expectations assumption above, the algorithm proceeds as follows:

1. **Initial State:** Start with the population distribution L_0 from observed data. Since the observed data does not perfectly match the model's stationary distribution, L_0 represents a non-steady-state initial condition.

2. **New Steady State:** Solve for the new equilibrium value functions (V_{SS}, W^c_{SS}) and decision rules (Λ_{SS}, π_{SS}) under the new parameters Θ_{new}.

3. **Simulation Loop:** For each generation $t = 1, \ldots, T$:

   - **Migration:** Young adults move according to the new steady-state transition matrix Λ_{SS}.

   $$L_t^{young}(e, j', k') = \sum_j L_{t-1}^{child}(e, j) \cdot \Lambda_{SS}(e, j, j', k')$$

   - **Market Clearing:** Housing prices p_t and wages Y_t adjust to clear local markets. Housing demand comes from both the young generation L_t^{young} and the old generation L_t^{old} (survivors from $t-1$).

   $$p_t(j) = \frac{1}{H_j} \left(\sum_{e,k} L_t^{young}(e, j, k) \sum_n P_t(n|e, j, k) \alpha(e, n) I_t(n) + \sum_e L_t^{old}(e, j) \alpha(e, 0) Y_t^o(e, j) \right)$$

   - **Population Evolution:** New children are born and educated, forming the next generation. We assume households are formed by couples, so the number of new children is scaled by 1/2. The education transition probability $\pi_{SS}(e \to e'|j, n)$ represents the probability that parents of type e in location j choose education e' for their n-th child. This decision maximizes the single-child value $W^{single}(e, j, n)$ given the steady-state continuation value $V^y_{SS}(e', j)$.

   $$L_t^{child}(e', j') = \sum_e \sum_{k'} \sum_{n'} \frac{\gamma_t(e, n', j', k') \cdot n'}{2} \pi_{SS}(e \to e'|j, n) L_t^{young}(e, j', k')$$

G Pronatalist Policies and Fertility: Empirical Estimates and Model Counterfactuals

This appendix provides a systematic comparison of estimated fertility effects across the pronatalist policy literature. Table A.4 compiles reduced-form estimates from quasi-experimental studies of cash transfers, childcare expansions, parental leave reforms, housing policies, and education cost interventions. Effect magnitudes are converted to percent changes where possible. Table A.5 summarizes counterfactual experiments from structural models of fertility, organized by mechanism. These tables complement the discussion in Section 1 and provide benchmarks for the policy experiments in Section 7.

Table A.4: Selected empirical estimates of pronatalist policies

Policy type	Paper	Country	Effect on fertility	Outcome	Identification
Panel A. Cash transfers and child allowances					
Family allowances	Gauthier and Hatzius (1997)	22 OECD	+0.07 TFR for +25% benefits (≈+4%)	TFR level	Cross-country panel
Baby bonus	Milligan (2005)	Canada	≈+12% avg; up to +25% at max	Births	Policy variation
Birth grant	González (2013)	Spain	≈+5–6% births	Births	RD
Child allowance	Cohen et al. (2013)	Israel	+1.05 pp (≈+7.5%)	Birth prob.	Benefit var.
Maternity capital	Slonimczyk and Yurko (2014)	Russia	+0.15 children (≈+12.7%)	Completed fert.	Structural
Family 500+	Bokun (2024)	Poland	+1.5 pp birth prob. (≈+27%)	Birth prob.	Eligibility
Panel B. Childcare subsidies and public provision					
Childcare slots	Rindfuss et al. (2010)	Norway	+0.5–0.7 children at 60% coverage	Completed fert.	Simulation

Continued on next page

Policy type	Paper	Country	Effect on fertility	Outcome	Identification
Childcare rollout	Bauernschuster et al. (2016)	Germany	+0.28% per 1 pp coverage	Births	Rollout DiD
Childcare avail.	Kravdal (1996)	Norway	≤+0.05 completed fert. for +20 pp	Completed fert.	Muni. var.
Panel C. Parental leave policies					
Leave extension	Lalive and Zweimüller (2009)	Austria	+15% prob. within 3 yrs; +3.0–3.5 pp in 10 yrs	Birth prob.	RD
Earnings-related	Raute (2019)	Germany	+22–23% for tertiary-educated	Birth prob.	Benefit DiD
Leave length	Dahl et al. (2016)	Norway	≈0% completed fertility	Completed fert.	Reform DiD
Speed premium	Hoem (1993)	Sweden	Large timing (tempo) response	Birth hazard	Reform
Panel D. Housing policies					
Housing credit	Fazio et al. (2024)	Brazil	+3.8% prob.; +3.2% number children	Birth prob.	Lottery
Housing wealth	Ang et al. (2024)	China	+7.3% fertility prob.	Birth prob.	RD/IV
House prices	Dettling and Kearney (2014)	USA	+5% (owners), −2.4% (non-owners)	Birth rate	MSA panel

Continued on next page

Policy type	Paper	Country	Effect on fertility	Outcome	Identification
Panel E. Education cost policies					
Education subsidy	Keng et al. (2011)	Taiwan	−17% when subsidy reduced	Birth prob.	Policy DiD
Panel F. Korea-specific evidence					
Baby bonus	Kim (2024)	S. Korea	+4.7% TFR from transfers	% TFR	Event-study

Notes: Effect magnitudes converted to percent changes where possible. Outcome: TFR = total fertility rate level; % TFR = percent change in TFR; Birth prob. = probability of birth; Completed fert. = children ever born; Birth rate = births per population. Cohen et al. (2013) estimate is for households below the poverty line (1.05 pp on a 13.92% baseline). Dettling and Kearney (2014) effects are per $10,000 house-price increase; "non-owners" is the original paper's terminology. Keng et al. (2011) estimate applies to public-sector households. Kim (2024) reports a counterfactual: the 2015 TFR would have been 4.7% lower absent municipal transfers. Kravdal (1996) reports the 0.05 estimate as an upper bound.

Table A.5: Structural Analyses of Fertility: Model-Based Counterfactuals

Paper	Model	Mechanism	Country	Counterfactual	Effect
Kim et al. (2024)	OLG/dynastic	Status externalities	S. Korea	Eliminate education arms race	+28% TFR
Fehr and Ujhelyiova (2013)	OLG + labor	Cash vs in-kind	Germany	Increase childcare (tax-financed)	TFR 1.46 to 1.88 (+29%); cash-to-in-kind shift: +4%
Bick (2016)	Life-cycle	Childcare costs	Germany	Childcare reforms	≈0% to −2% fertility
Guner et al. (2024)	Life-cycle	Labor frictions	Spain	Reduce duality + expand childcare	Completed fert. 1.60→1.74 (no college) and 1.60→1.96 (college)
Sommer (2016)	Life-cycle	Earnings risk	USA	Higher earnings risk	−14% births; access to IVF offsets by +5%
Da Rocha and Fuster (2006)	Life-cycle	Unemployment	Spain/US	Lower unemployment	+19% TFR
Couillard (2025)	Housing model	Housing costs	USA	Undo post-1990 rent increases	+11% births
Doepke and Kindermann (2019)	Bargaining	Childcare burden	GGP countries	Reduce mothers' burden	Policies scaled to +0.1 TFR; cost per couple varies (18k–95k EUR)

Notes: Effect is percent change in TFR or completed fertility relative to baseline unless otherwise noted. Guner et al. (2024) reports levels by education group. Doepke and Kindermann (2019) hold the fertility increase fixed at 0.1 and compare the per-couple cost across policy instruments. Most papers report additional experiments.

Table A.6: Sensitivity to Dispersion Parameters: Counterfactual Experiments

	Baseline Parameter		σ_N (0.5)		σ_N (1.2)		σ_{JK} (0.10)		σ_{JK} (0.25)		σ_E (0.15)		σ_E (0.6)	
Scenario	TFR	%Δ	TFR	%Δ	TFR	%Δ	TFR	%Δ	TFR	%Δ	TFR	%Δ	TFR	%Δ
Baseline	1.513	–	1.498	–	1.493	–	1.489	–	1.519	–	1.357	–	1.806	–
Education Cost (−20%)														
Seoul only	1.519	0.4	1.538	2.6	1.508	1.1	1.515	1.7	1.542	1.5	1.389	2.4	1.815	0.5
All regions	1.755	14.8	1.833	20.2	1.672	11.4	1.752	16.2	1.776	15.6	1.625	18.0	2.037	12.1
Subsidies to Parents														
Childless tax (20%)	1.630	7.5	1.675	11.2	1.587	6.1	1.627	8.9	1.651	8.3	1.514	10.9	1.898	5.0
Universal tax	1.536	1.5	1.549	3.3	1.524	2.1	1.533	2.9	1.558	2.5	1.413	4.1	1.826	1.1
Edu Reform + Subsidy														
Edu + childless	1.855	20.3	1.947	26.2	1.750	15.9	1.852	21.8	1.872	20.9	1.743	25.1	2.099	15.0
Edu + universal	1.789	16.7	1.870	22.2	1.699	13.0	1.786	18.2	1.807	17.4	1.670	20.8	2.053	12.8
Housing Supply (+100%)														
Seoul only	1.466	-3.1	1.467	-2.1	1.473	-1.4	1.461	-1.9	1.500	-1.2	1.330	-2.0	1.780	-1.4
All regions	1.499	-0.9	1.510	0.8	1.495	0.2	1.495	0.4	1.525	0.4	1.364	0.5	1.808	0.1
Commuting (−20% CEV)														
Seoul only	1.453	-4.1	1.449	-3.3	1.464	-1.9	1.447	-2.9	1.492	-1.8	1.316	-3.0	1.770	-2.0
All regions	1.494	-1.3	1.499	0.1	1.493	0.0	1.491	0.1	1.520	0.1	1.358	0.1	1.807	0.1
Amenity (+20% CEV)														
Seoul only	1.439	-5.1	1.431	-4.5	1.454	-2.7	1.431	-4.0	1.482	-2.5	1.302	-4.1	1.756	-2.8
All regions	1.495	-1.2	1.500	0.1	1.494	0.1	1.491	0.1	1.520	0.1	1.358	0.1	1.807	0.1
Productivity (TFP +100%)														
Seoul only	1.409	-7.2	1.393	-7.3	1.433	-4.1	1.405	-5.9	1.453	-4.4	1.272	-6.5	1.727	-4.4
All regions	1.493	-1.4	1.493	-0.3	1.494	0.1	1.489	-0.0	1.518	-0.0	1.356	-0.1	1.808	0.1

H Additional Transition Results

Figure A.2: Transition Dynamics: Housing and Commuting

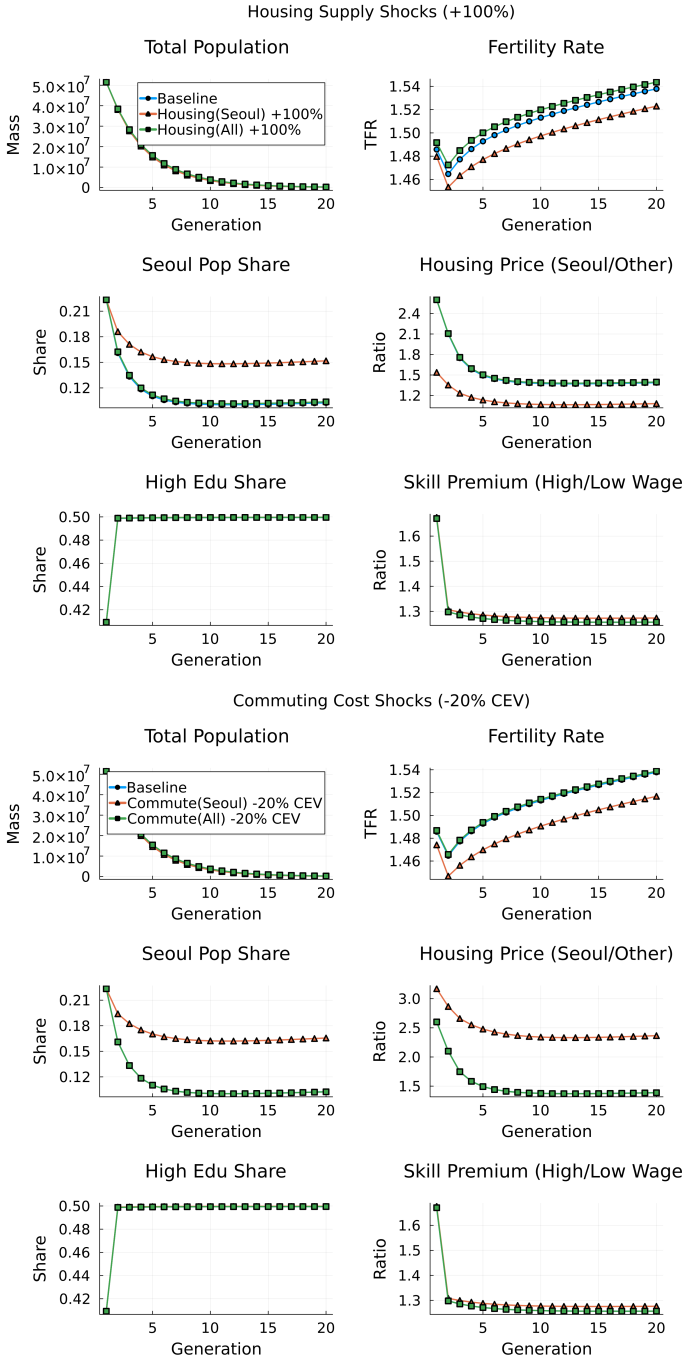


Figure A.3: Transition Dynamics: Amenity and Productivity

