Where to Build New Cities: Optimal Urban Land Supply and Korea's New Town Projects*

Seungyub Han† Sunham Kim‡

June 26, 2026

Abstract

Where should governments supply new urban land? We study this question in a dynamic spatial growth model with a government that allocates land across regions under a budget constraint, where procurement cost scales with the local land price. In the short run, the welfare gain per dollar of procurement cost is equal across regions, and location is irrelevant at the margin. Permanently expanding the urban land stock moves the economy to a new steady state, as capital accumulates, workers migrate, and prices adjust over decades. This long-run gain per dollar no longer scales with the local land price. It varies across regions with the reallocation and accumulation of capital and labor that the new land induces, and the welfare return per dollar of procurement cost, not the local land price, ranks regions for new land. We take the model to South Korea's Second New Town Project, a large-scale urban land supply that converted roughly 5,000 hectares around Seoul into nine serviced cities and delivered around 610,000 housing units at a cost above 10% of 2004 GDP. The project raised long-run aggregate welfare by 0.28%. In our counterfactual, the same budget supplied to a lower-priced region outside the Seoul Metropolitan Area raises welfare by 0.35%. Our results suggest that when the instrument for correcting spatial misallocation is costly, where to deploy it depends on the long-run gains from adjustment relative to their cost, which may not necessarily be the most productive regions.

Keywords: Active land supply, Housing supply, Spatial misallocation, New Town Projects

JEL Classification: E24, H54, R11, R13, R14, R52, R58

*We deeply appreciate the guidance and advice given by Lee Ohanian and Mario Crucini. We appreciate helpful comments from Jan K. Brueckner and Anthony M. Yezer. This research is supported by the Rosalinde and Arthur Gilbert Research Grant (UCLA Ziman Center for Real Estate, 2023) and the Krannert Doctoral Research Funds (Purdue University, 2023). Any errors are our own.

†Louisiana State University (email: hansy1124@lsu.edu)

‡Korea Development Institute (email: sunham.kim@kdi.re.kr)

1 Introduction

Scarce urban land constrains housing supply, and constrained housing supply limits the inflow of workers and capital into productive regions. The resulting spatial misallocation imposes substantial losses in aggregate output and welfare (Glaeser, 2014). The gains from expanding effective urban land through deregulation, which intensifies the use of existing land in the most productive cities, are studied extensively (Hsieh and Moretti, 2019; Herkenhoff et al., 2018; Greaney, 2026). Less is known about active land supply: where to convert rural land into new urban acres when its cost, tied to local land prices, rises across space as steeply as the benefit.

Public new-city programs, which convert rural land into urban use, housed growing populations across the mid-twentieth century and relieve housing shortages in congested cities today. England's postwar New Towns produced more than 300,000 housing units across about 30 towns (United Kingdom Parliament, 1946), and France's *Villes nouvelles* program developed nine planned towns as part of a national policy linking housing production with urban planning (Korsak and Pernelle, 2004).[1] More recently, the Australian state of Victoria has planned some 180,000 homes on greenfield land around Melbourne (Department of Transport and Planning, Victoria, 2024), and the United Kingdom has revived the instrument under a renewed new-towns agenda (Ministry of Housing, Communities and Local Government, 2024). South Korea's Second New Town Project, a large case of active land supply, converted roughly 5,000 hectares around Seoul into nine serviced cities and delivered around 610,000 housing units over the 2000s. Each parcel was bought at its local rural land price, then serviced with roads, utilities, and public infrastructure before any household could move in. The bill came to 95 trillion KRW, above 10% of South Korea's 2004 GDP.

In this paper, we ask where a government should supply new urban land. Once supplied, urban land cannot revert to rural use. Permanently expanding the stock of land, a fixed factor, moves the economy to a new steady state over decades, as capital accumulates and workers migrate across regions. Land supply is also costly. Beyond servicing it with infrastructure, the government must buy the land, at the local price.[2] The benefit accrues mainly to the region that receives the land, and to the aggregate economy through the adjustment that follows. Together, these features make active land supply a problem of spatial investment under a resource constraint.

[1] Japan is another case. Its Tama New Town, begun in 1965 west of Tokyo, was planned for 342,000 residents on 2,900 hectares (Urban Renaissance Agency, 2007). It was one of several towns developed under the New Residential Town Development Act of 1963.

[2] Whether the government acquires the land directly or through private developers, the cost still reflects the local land price.

We study this problem in the dynamic spatial growth model of Herkenhoff et al. (2018) (hereafter HOP), introducing a government that supplies land in each region under a fixed budget.[3] The welfare benefit of new land has two horizons. In the short run, before the economy adjusts, the benefit equals the discounted value of the local rent the new land earns. Procurement cost is proportional to that same rent, so the return per dollar is equal across regions, leaving the government indifferent about where to build. In the long run, the added land expands the region's production frontier, raises the marginal product of capital, and draws in a larger capital stock. We decompose this long-run gain into three channels: the complementarity between land and capital in production, the substitution between land and capital in housing, and labor adjustment, combining reallocation across regions with the aggregate labor-supply response. Through these channels the long-run gain no longer scales with the local land price, and the return per dollar of procurement cost ranks regions for new land.

South Korea and its Second New Town Project (2004) make a natural setting for this question. The country supplied urban land through public programs, and the Second New Town Project is the largest of them. The government set how much land each site received, and procurement cost and land supply are reported at the site level. This information enables us to quantify how much the same budget would have supplied in other regions. We calibrate the model to fifteen Korean provinces over 2004–2019. Region-specific productivity, amenity, and land-use restriction are recovered from the model's equilibrium conditions following HOP, given the urban land stock from official cadastral statistics. The calibrated model matches untargeted features of the data, both the fundamentals it infers and the prices it predicts. The inferred land-use restriction correlates with the floor area and building coverage ratios, and the inferred amenities with a survey-based index of regional quality of life. Model-implied land rents track observed transaction prices closely in level and fit, which matters because land prices drive the cost in our framework and the welfare comparisons depend on their ranking across regions.

We use the calibrated model to evaluate the Second New Town Project, which supplied new land to the Seoul Metropolitan Area (SMA). Relative to the no-NTP counterfactual, this allocation raised long-run aggregate welfare by 0.28% in consumption-equivalent terms and output by 0.23%. The welfare-maximizing allocation under the same budget, by contrast, would have concentrated new land in Daejeon, outside the SMA, raising welfare by 0.35%. Daejeon ranks below Seoul in raw marginal welfare, yet delivers the higher return per procurement dollar, because its land price is far

[3] We build on HOP because its dynamic structure captures the economy's adjustment across steady states, over which the long-run gains from land supply accumulate.

lower. Part of the welfare benefit, the discounted rent on the new land, is proportional to that same price, so it cancels in the ratio. The rest of the benefit comes from the long-run adjustment. Of its three channels, the production-side complementarity and the housing-side substitution are nearly identical across the two allocations. Labor adjustment accounts for the entire difference. Additional serviced land in Daejeon relaxes its housing constraint and raises the return to working there, drawing in workers and expanding aggregate labor supply. The same budget in the dense, expensive SMA supplies far less serviced land, so it loosens housing constraints little and the labor-supply response is weak. The actual project thus captured about 80% of the welfare available under the same-budget optimum, while deepening the SMA's concentration of population and economic activity, a long-standing concern of Korean spatial policy.

The welfare return per procurement dollar varies systematically with a region's fundamentals. It depends on their full joint distribution, and the non-linear general-equilibrium response rules out a closed form, but the calibrated economy reveals some clear associations. The return is higher in regions with greater amenities. A more attractive region draws a larger labor-supply response to new land and, through the complementarity between labor and capital, a larger long-run expansion in capital and output. The return is lower where the existing stock of effective urban land is large. Productivity has two offsetting effects. More productive regions yield larger gross welfare gains from new land, but they also command higher land prices, and the higher procurement cost offsets much of the gain, leaving its net association with the return unclear. In sum, a region's productivity alone does not make it the most cost-effective place for new land.

Our analysis builds directly on the spatial-misallocation literature on land-use regulation. Glaeser (2014) and Furman (2015) document the link between binding land-use restrictions, inelastic housing supply, and slowing regional convergence in the United States. Hsieh and Moretti (2019) and Herkenhoff et al. (2018) formalize the aggregate consequences in spatial general-equilibrium settings, showing that stringent regulation in productive cities lowers aggregate output substantially. Babalievsky et al. (2026) extend this framework to commercial real estate, and Albouy and Ehrlich (2018) estimate the social cost of land-use restrictions via a metro-level housing cost function. Our framework shares this literature's emphasis on spatial misallocation and its structural treatment of regional fundamentals, but it shifts the policy margin from the regulatory wedge to the urban land stock, financed by an explicit budget. This shift is essential for our central object, the procurement-cost-normalized welfare return: deregulation in the existing literature is costless from the planner's perspective, whereas active land supply trades real fiscal resources for additional land. The optimal allocation rule we derive therefore takes the form of a welfare gain per dollar of procurement cost, which has no direct analogue in the deregulation literature.

A related body of work studies place-based policies and their regional incidence. Kline and Moretti (2014) document large persistent gains from the Tennessee Valley Authority, a big-push program; Busso et al. (2013) provide a canonical incidence-and-efficiency analysis of US Empowerment Zones; and Neumark and Simpson (2015) survey the broader literature. Donaldson and Hornbeck (2016) evaluate the welfare consequences of large public infrastructure (US railroads) using a market-access spatial framework. Moretti (2010) measures local employment multipliers from tradable-sector jobs, and Koster and van Ommeren (2019) document the housing-market response to place-based investment. We also connect to the regional decline literature documenting the slowdown in inter-regional income convergence and skill sorting (Moretti, 2013; Ganong and Shoag, 2017; Diamond, 2016). Relative to these works, we characterize an optimal cross-regional allocation rule and quantify its national-level general-equilibrium effects, rather than estimating the local incidence of an existing policy. In our application, we additionally bring the regional-decline lens to examine how outcomes in declining peripheral regions respond to active urban land supply policy.

Our framework belongs to the broader quantitative spatial general-equilibrium literature. The static spatial-equilibrium foundation traces to Roback (1982); quantitative spatial general equilibrium methodology has been developed by Allen and Arkolakis (2014) and surveyed by Redding and Rossi-Hansberg (2017). Dynamic spatial general equilibrium with frictional labor mobility is developed by Caliendo et al. (2019), and Desmet and Rossi-Hansberg (2014) integrate endogenous innovation into a spatial growth model. We differ from these papers in two ways. First, we embed a neoclassical growth structure: capital accumulation responds to land supply through marginal-product complementarity, and the economy converges to a new long-run steady state. While we do not endogenize productivity growth, our framework lets us trace how a permanent expansion of urban land in any region propagates through capital accumulation, labor migration, and prices into long-run welfare changes within an otherwise standard growth model. Second, we treat urban land supply as an allocation problem for the government, the choice of where to spend a fixed public budget when procurement cost and long-run welfare return vary across regions.

The remainder of the paper proceeds as follows. Section 2 sets up the spatial growth model and the government's land-supply problem. Section 3 develops the welfare-derivative decomposition, the procurement-cost-normalized welfare return, and the optimal allocation rule. Section 4 describes the institutional context of South Korea and its New Town Projects. Section 5 calibrates the model to Korean regional data and validates the fit against untargeted moments. Section 6 reports the counterfactual experiments. Section 7 concludes.

2 A Spatial Model with Active Land Supply

We develop a multi-region spatial growth model in which a benevolent government chooses how much new urban land to supply in each region. The framework extends Herkenhoff et al. (2018) (hereafter HOP) by introducing an explicit government that optimizes the long-run supply of urban land across regions, subject to a real-resource budget and a region-specific development cost.

2.1 Environment

Time is indexed by $t = 0, 1, 2, \ldots$ and regions by $j \in \{1, \ldots, J\}$. Each region j is characterized by three exogenous regional fundamentals—total factor productivity A_j, amenity a_j, and land-use regulation $\alpha_j \in (0, 1]$—together with an urban land endowment $x_{jt} > 0$ owned by the representative household. A benevolent government can make a one-time decision to expand the urban land endowment $\{x_{jt}\}$ at a real-resource cost; we defer the timing, pricing, and budget details of this policy to Section 2.3. Until then, we take $\{x_{jt}\}$ as given.

Households. A single representative household owns the aggregate capital stock and all regional land endowments, supplies labor across regions, consumes final goods and housing services, and saves. Let n_{jt} denote employment in region j. The household solves the following problem:

$$\max \quad \sum_{t=0}^{\infty} \beta^t \left[\ln c_t - \sum_{j=1}^{J} \frac{n_{jt}^{1+\frac{1}{\gamma}}}{1+\frac{1}{\gamma}} + \sum_{j=1}^{J} a_j\, n_{jt} \right] \tag{1}$$

$$\text{s.t.} \quad c_t + i_t + \sum_{j=1}^{J} p_{jt}\, h_{jt} \;=\; r_t\, k_t + \sum_{j=1}^{J} \big(w_{jt}\, n_{jt} + q_{jt}\, x_{jt} + \pi_{yjt} + \pi_{hjt} \big), \tag{2}$$

$$i_t = k_{t+1} - (1-\delta)\, k_t, \qquad k_t = \sum_{j=1}^{J} \big(k_{yjt} + k_{hjt} \big), \tag{3}$$

$$x_{jt} = x_{yjt} + x_{hjt}, \qquad \forall j,\, t, \tag{4}$$

$$h_{jt} \geq n_{jt}, \qquad \forall j,\, t. \tag{5}$$

Here, $\beta \in (0,1)$ is the discount factor, c_t is aggregate consumption, and $\gamma > 0$ governs the curvature of region-specific labor disutility. The convex labor-disutility term provides a reduced-form regional congestion force, in addition to the congestion generated by housing and land market clearing. The term $a_j n_{jt}$ captures the amenity value associated with employment in region j. The budget constraint (2) equates expenditures and income. On the expenditure side, the household purchases consumption c_t, makes aggregate capital investment i_t, and pays p_{jt} per unit of housing services h_{jt} in each region. On the income side, it rents the aggregate capital stock k_t to firms at the economy-wide rate r_t—equalized across regions and sectors by free capital mobility—earns wage income $w_{jt}\, n_{jt}$ in each region, collects land-rental income $q_{jt}\, x_{jt}$ from its regional land endowment, and receives profits π_{yjt} and π_{hjt} distributed by the region-j final-goods and housing firms, respectively.

The remaining constraints describe capital, land, and housing. Equation (3) is the standard capital law of motion with depreciation rate $\delta \in (0,1)$, together with the accounting identity that aggregate capital k_t equals the sum of capital rented to final-goods firms (k_{yjt}) and housing firms (k_{hjt}) across regions. Equation (4) splits the region-j land endowment between final-goods production, x_{yjt}, and housing production, x_{hjt}. The housing-feasibility constraint (5) requires each worker in region j to occupy at least one unit of housing services. Because $p_{jt} > 0$ in equilibrium and housing provides no direct utility, the constraint always binds: $h_{jt} = n_{jt}$.

Final goods. A representative firm in each region produces a tradable final good:

$$y_{jt} = A_j\, \widehat{A}(\bar{y}_{jt})\, \bar{y}_{jt}, \qquad \bar{y}_{jt} = k_{yjt}^{\theta}\, n_{jt}^{\chi}\, (\alpha_j\, x_{yjt})^{\varphi}, \tag{6}$$

with $\varphi \equiv 1 - \theta - \chi > 0$. The land-use regulation $\alpha_j \in (0,1]$ enters as an effective-use multiplier: supplying x_{yjt} units of raw land yields only $\alpha_j x_{yjt}$ units of productive land input, capturing real-world restrictions such as floor area ratios and building coverage ratios that cap how intensively a given plot can be used. The firm takes $\widehat{A}(\bar{y}) = \bar{y}^{\lambda}$ as given, so that $\lambda \geq 0$ parameterizes an external agglomeration force. At $\lambda = 0$ there is no externality and the technology has constant returns to scale; at $\lambda > 0$, $\widehat{A}$ generates a positive production externality and the aggregate technology exhibits increasing returns to scale. The firm solves

$$\pi_{yjt} = \max_{k_{yjt},\, n_{jt},\, x_{yjt}} \quad y_{jt} - r_t k_{yjt} - w_{jt} n_{jt} - q_{jt} x_{yjt}. \tag{7}$$

Housing. A representative rental housing firm in each region produces with

$$h_{jt} = k_{hjt}^{\xi} \left(\alpha_j \, x_{hjt}\right)^{1-\xi}, \tag{8}$$

where $1 - \xi \in (0, 1)$ is the land-input share in housing production, and ξ is the corresponding capital share. The firm solves

$$\pi_{hjt} = \max_{k_{hjt}, x_{hjt}} \quad p_{jt} h_{jt} - r_t k_{hjt} - q_{jt} x_{hjt}. \tag{9}$$

2.2 Competitive Equilibrium

Given regional fundamentals $\{A_j, a_j, \alpha_j\}_{j=1}^{J}$ and a regional land endowment $\{x_j\}_{j=1}^{J}$, a competitive equilibrium is a sequence of quantities $\{c_t, k_{t+1}, \{n_{jt}, h_{jt}, k_{yjt}, k_{hjt}, x_{yjt}, x_{hjt}\}_{j=1}^{J}\}_{t \geq 0}$ and prices $\{r_t, \{w_{jt}, p_{jt}, q_{jt}\}_{j=1}^{J}\}_{t \geq 0}$ such that (i) the household optimizes, (ii) each firm optimizes taking $\widehat{A}$ as given, (iii) final-goods, labor, capital, land, and housing markets clear in every region and period, and (iv) housing feasibility $h_{jt} \geq n_{jt}$ holds—enough housing is supplied for every worker in region j.

2.3 Government and the Land Supply Policy

Prior to the policy, the economy is assumed to be in a long-run steady state with regional land endowments $\{x_{j,0}\}_{j=1}^{J}$, $x_{j,0} > 0$, and pre-policy steady-state land rental rates $\{q_{j,0}\}_{j=1}^{J}$. At $t = 0$, observing all regional fundamentals, the government makes a one-time decision to permanently augment region j's urban land endowment by supplying $\Delta x_j \geq 0$ additional units, financed by a real-resource budget B. The supplied land is transferred to the representative household's regional land endowment, so that $x_{jt} = x_{j,0} + \Delta x_j$ for all $t \geq 0$, and the economy transitions to a new long-run steady state.

Under this set-up, the government solves the following problem:

$$\max_{\{\Delta x_j\}_{j=1}^{J}} \quad W_{\text{ss}}(x_{1,0} + \Delta x_1, \ldots, x_{J,0} + \Delta x_J) \tag{10}$$

$$\text{s.t.} \quad \sum_{j=1}^{J} \frac{\kappa \, q_{j,0}}{1-\beta} \Delta x_j \; \leq \; B, \qquad \Delta x_j \geq 0,$$

where $W_{\text{ss}}(x)$ is the steady-state lifetime utility of the representative household evaluated at the stationary competitive equilibrium associated with the land vector x—i.e., the present value of flow utility at the new steady state.[4]

[4] The full transition dynamics from the pre-policy to the post-policy steady state are also of interest,

Under this budget constraint, the per-unit *procurement cost* (or, equivalently, *development cost*) of supplying urban land in region j is $\kappa\, q_{j,0}/(1-\beta)$, which decomposes into two pieces. The first, $q_{j,0}/(1-\beta)$, is the pre-policy *asset value* of a unit of region-j land—the present-discounted value of pre-policy flow rents at the steady state. The second, $\kappa > 0$, is a real-resource scaling factor that maps the pure asset value into the actual real-resource cost the government incurs in converting one unit of raw land into usable urban land. Three regimes are economically meaningful. $\kappa = 1$ corresponds to procurement at exactly the competitive-equilibrium land price—the pure-asset-value benchmark. $\kappa > 1$ captures situations in which actual development costs exceed the asset value, reflecting complementary infrastructure, displacement, regulatory, and administrative outlays on top of land acquisition. $\kappa < 1$ corresponds to situations in which the government can develop peripheral parcels within a region at a discount relative to its central land price. We treat κ as a common, region-invariant scalar throughout the theoretical analysis.

Tying the per-unit cost to the asset-value anchor $q_{j,0}/(1-\beta)$ is consistent with the real-world economics of urban land development. In already-developed metropolitan areas, where land rents $q_{j,0}$ are typically high, the easiest parcels have often already been exhausted, leaving sites that are scarcer and more costly to acquire and convert to urban use. More generally, undeveloped land near large metropolitan areas commands a higher market value than comparable land near smaller cities, raising procurement and development costs. It is therefore natural, as a reduced-form benchmark, to assume that regions with higher land rents $q_{j,0}$ also face proportionally higher per-unit costs of supplying new urban land. We further assume that $q_{j,0}$ in the government budget constraint is evaluated at the pre-policy steady state and is not affected by the government's subsequent policy announcement. This treatment is consistent with a common feature of public land-acquisition frameworks, in which procurement and compensation costs are anchored to prevailing market values at a specified valuation date, rather than to land values subsequently altered by the project itself. Korea's Second New Town Project provides a concrete example: although the broader program was announced in advance, the specific development sites remained confidential until their official designation, limiting the extent to which project-specific information was capitalized into pre-acquisition land prices.

---

but as a first pass we target the new steady-state welfare and leave the full transition analysis to future research. This is a reasonable benchmark for large permanent land supply programs such as the Second NTP, whose reallocative effects unfold over decades. In particular, the cost of the policy is a one-time real-resource expenditure while the output and welfare gains are permanent.

3 Theoretical Analysis

The government's problem (10) is fundamentally a question about which regions deliver the greatest long-run welfare gain per dollar of procurement cost. Solving it directly requires computing a new post-policy steady state for each candidate allocation—tractable only numerically, as in our quantitative analysis of Section 6.2. Analytical progress, however, is available at the margin: the welfare gain per dollar spent on region-j land supply, evaluated at the pre-policy steady state, is governed by the derivative $dW_{ss}/d\Delta x_j$ and its normalization by the per-unit procurement cost $\kappa\, q_{j,0}/(1-\beta)$. This section studies that object—decomposing it into economically interpretable channels, characterizing how it varies across regions, and deriving the implications for the optimal allocation.

Throughout, we evaluate how a permanent urban land supply affects the long-run steady state of the decentralized competitive equilibrium, taking first-order derivatives at the pre-policy steady state. We assume that the stationary competitive equilibrium selected at the pre-policy state is locally unique and differentiable in the regional land-endowment vector x in a neighborhood of x_0.

3.1 The Welfare Derivative Decomposition

Theorem 1 (Welfare derivative decomposition at the pre-policy competitive equilibrium)**.** *At the pre-policy steady state (i.e., at $\Delta x = 0$), let $q_{j,0}$ denote the market land rent, c_0 the aggregate consumption, and $u_c \equiv u_c(c_0)$ the marginal utility of consumption. The welfare derivative with respect to the region-j land supply admits the decomposition*

$$\frac{dW_{ss}}{d\Delta x_j} = \underbrace{\frac{u_c}{1-\beta} q_{j,0}}_{\text{direct land channel}} + \underbrace{\frac{u_c}{\beta} \frac{dK_{ss}}{d\Delta x_j}}_{\text{capital accumulation channel}} + \underbrace{\frac{\lambda}{1+\lambda} \cdot \frac{u_c}{1-\beta} \cdot \frac{dY_{ss}}{d\Delta x_j}}_{\text{externality channel}}, \tag{11}$$

where all derivatives are evaluated at $\Delta x = 0$.

Proof. See Appendix B.1. □

The three components of (11) admit an intuitive reading. The *direct land channel* captures the welfare gain when all other production factors—especially aggregate capital—are held at their pre-policy steady-state values and only the urban land endowment in region j rises: the household earns an additional $q_{j,0}$ per period in land-rental income, capitalized at the present-value factor $(1-\beta)^{-1}$. Here $q_{j,0}$ is the private marginal product of land in region j—how much more goods or housing the economy can produce with one more unit of land at decentralized-equilibrium prices, excluding any agglomeration spillover. This term scales with $q_{j,0}$ and so introduces an apparent cross-regional heterogeneity that will, as we show below, vanish once the welfare derivative is normalized by a procurement cost that itself scales with $q_{j,0}$.

The *capital accumulation channel* is the central long-run gain mechanism of the policy. When the government supplies new urban land in region j, the land is absorbed by both the production sector and the housing sector of region j, and firms in each sector adjust their capital demand. Production-side capital tends to rise through complementarity (more land raises the marginal product of capital, drawing in capital), while housing-side capital substitutes against the additional land at the binding housing constraint $h_\ell = n_\ell$. At the same time, the new land in region j also raises regional labor demand through complementarity and triggers cross-region labor reallocation toward j, which feeds back into both the production-side capital demand and the housing-side capital required to house the migrating workers. The household responds to the net effect by adjusting the steady-state aggregate capital stock until the Euler condition $\beta(r + 1 - \delta) = 1$ is restored, and the resulting extra consumption contributes to welfare through u_c/β per unit of induced capital.[5] In other words, Δx_j changes the production frontier of the whole economy while causing reallocation of production factors, all of which is summarized by the state variable, aggregate capital K_{ss}. The induced response $dK_{\text{ss}}/d\Delta x_j$ has no reason to scale with $q_{j,0}$ and instead reflects the joint distribution of regional fundamentals across all regions; as we develop below, it is therefore the source of cross-regional heterogeneity in the land-supply impact that survives even the procurement-cost normalization. We unpack the response in the next subsection.

The *externality channel* captures the social output gain that the competitive market fails to price: at the decentralized equilibrium, firms do not internalize the agglomeration spillover, so the market rent $q_{j,0}$ understates the social return by exactly the $\lambda/(1 + \lambda)$ fraction of the induced output response $dY_{\text{ss}}/d\Delta x_j$. Although we have isolated this term separately for accounting clarity, the agglomeration force itself ultimately operates through the same general-equilibrium adjustment as the capital channel: at the competitive equilibrium $dY_{\text{ss}}/d\Delta x_j$ depends on the same regional capital response $dK_{\text{ss}}/d\Delta x_j$, an equivalence we record explicitly in Appendix C.1. The takeaway is unified: allocating new urban land to the region with the largest aggregate capital response $dK_{\text{ss}}/d\Delta x_j$ delivers a higher long-run welfare gain than the same allocation elsewhere—directly through the capital accumulation channel and

[5] The coefficient u_c/β admits a clean steady-state interpretation. One additional unit of steady-state capital generates gross flow MPK of r but requires extra depreciation δ each period to be maintained, yielding net flow consumption of $r - \delta$. Capitalizing this perpetuity at $1/(1 - \beta)$ and using the Euler condition $r - \delta = (1 - \beta)/\beta$ delivers the welfare contribution $u_c \cdot (r - \delta)/(1 - \beta) = u_c/\beta$ per unit of induced capital.

indirectly through the externality channel (agglomeration amplification).

3.2 Capital Response Decomposition

Because the aggregate-capital response $dK_{ss}/d\Delta x_j$ summarizes the equilibrium adjustment to a regional urban land supply and is the central source of cross-regional heterogeneity in the welfare derivative (11), we now decompose it in terms of primitive regional allocations rather than reduced-form objects such as Y_{ss} or p_ℓ. This yields a cleaner economic reading because each channel maps to a distinct land- or labor-reallocation mechanism and exposes how the agglomeration externality λ enters each channel separately—directly amplifying production-sector responses while leaving housing-sector responses untouched.

Theorem 2 (Capital response decomposition in primitives). *At the steady-state with $\theta(1+\lambda) < 1$,*[6] *the derivative of aggregate steady-state capital with respect to the region-j land endowment admits the three-channel decomposition*

$$\frac{dK_{ss}}{d\Delta x_j} = \underbrace{\sum_\ell \Gamma_\ell^{xy} \frac{dx_{y\ell}}{d\Delta x_j}}_{\substack{\text{land–capital complementarity} \\ \text{(production)}}} - \underbrace{\sum_\ell \Gamma_\ell^{xh} \frac{dx_{h\ell}}{d\Delta x_j}}_{\substack{\text{land–capital substitution} \\ \text{(housing)}}} + \underbrace{\sum_\ell \Gamma_\ell^{\text{lab}} \frac{dn_\ell}{d\Delta x_j}}_{\text{labor adjustment}}, \tag{12}$$

with coefficients

$$\Gamma_\ell^{xy} \equiv \frac{k_{y\ell}}{x_{y\ell}} \cdot \frac{\varphi(1+\lambda)}{1-\theta(1+\lambda)}, \quad \Gamma_\ell^{xh} \equiv \frac{k_{h\ell}}{x_{h\ell}} \cdot \frac{1-\xi}{\xi}, \quad \Gamma_\ell^{\text{lab}} \equiv \frac{k_{y\ell}}{n_\ell} \cdot \frac{\chi(1+\lambda)}{1-\theta(1+\lambda)} + \frac{k_{h\ell}}{n_\ell} \cdot \frac{1}{\xi},$$

all evaluated at the pre-policy steady state.

Proof. See Appendix B.2. □

The three channels of (12) decode the capital response into distinct economic mechanisms. Given that the underlying relationships are entirely among equilibrium objects, it is generally difficult to attribute the capital response to any specific parameter or regional fundamental in closed form; the value of the decomposition is to identify *which* channels mediate the aggregate capital adjustment, conceptually, rather than to provide an analytical map from primitives to channel magnitudes.

[6] The condition $\theta(1+\lambda) < 1$ rules out increasing returns to capital alone. The product $\theta(1+\lambda)$ is the *social* elasticity of regional output with respect to capital—the private share θ scaled by the agglomeration multiplier $(1+\lambda)$ that captures the spillover one firm's capital generates for others. If it exceeded one, the agglomeration-amplified return to capital would drive unbounded capital accumulation and no interior steady state would exist. For empirically reasonable capital shares (θ around 0.3) and modest agglomeration spillovers, the condition is satisfied with substantial slack.

The *land–capital complementarity* channel (Γ_ℓ^{xy}) captures the direct effect of more land allocated to the production sector: production complementarity raises the marginal product of capital, and firms scale up their capital holdings to maintain MPK $= r$. The agglomeration externality enters this channel through the elasticity factor $(1+\lambda)/(1-\theta(1+\lambda))$, capturing the spillover loop within the production sector: additional capital raises regional output, which raises productivity through the spillover, which in turn draws in still more capital.

The *land–capital substitution* channel (Γ_ℓ^{xh}) captures the direct effect of more land allocated to the housing sector: with $h_\ell = n_\ell$ binding, the fixed housing-services requirement pins the isoquant, so additional land lets housing firms economize on capital, shrinking $k_{h\ell}$. Concretely, when land is abundant, the same number of housing units can be provided with lower-rise, less capital-intensive structures rather than tall buildings, reducing the capital required per unit of housing. The channel does not carry an agglomeration multiplier because the housing technology does not host the spillover.

The *labor adjustment* channel (Γ_ℓ^{lab}) packages two sub-effects of a regional labor shift: a production-sector contribution $(k_{y\ell}/n_\ell)\cdot\chi(1+\lambda)/(1-\theta(1+\lambda))$, reflecting labor–capital complementarity that scales $k_{y\ell}$ with n_ℓ and carries the same agglomeration multiplier as Γ_ℓ^{xy}; and a housing-sector contribution $(k_{h\ell}/n_\ell)\cdot(1/\xi)$, reflecting that more workers mechanically raise housing demand—and thus $k_{h\ell}$—through the binding feasibility constraint. The agglomeration multiplier appears only in the production-sector contribution; even so, this multiplier is only one of several routes through which λ shapes $dK_{\text{ss}}/d\Delta x_j$, since the endogenous responses $\{dx_{y\ell}, dx_{h\ell}, dn_\ell\}$ that multiply the Γ's also depend on λ through the same equilibrium system. Appendix C.2 records the equivalent elasticity form of (12).

A useful refinement of the labor adjustment channel separates the two margins of the labor response. Under the current utility specification, total labor $N \equiv \sum_\ell n_\ell$ is endogenous, and $\sum_\ell dn_\ell/d\Delta x_j$ need not vanish. Writing $s_\ell \equiv n_\ell/N$, the product rule applied to $n_\ell = s_\ell N$ gives $dn_\ell = s_\ell\, dN + N\, ds_\ell$. Defining the cross-region reallocation component $d\tilde{n}_\ell \equiv N\, ds_\ell$—which satisfies $\sum_\ell d\tilde{n}_\ell = N\sum_\ell ds_\ell = 0$ identically since $\sum_\ell s_\ell = 1$—the decomposition $dn_\ell = d\tilde{n}_\ell + s_\ell\, dN$ isolates pure labor reallocation at fixed aggregate N from the aggregate labor-supply margin. The labor adjustment channel of (12) splits correspondingly as

$$\sum_\ell \Gamma_\ell^{\text{lab}} \frac{dn_\ell}{d\Delta x_j} = \underbrace{\sum_\ell \Gamma_\ell^{\text{lab}} \frac{d\tilde{n}_\ell}{d\Delta x_j}}_{\text{cross-region reallocation}} + \underbrace{\left(\sum_\ell s_\ell \Gamma_\ell^{\text{lab}}\right) \frac{dN}{d\Delta x_j}}_{\text{aggregate labor supply}}. \tag{13}$$

The first term is pure migration at fixed aggregate labor; the second captures the intensive-margin labor-supply response that arises from the Frisch-elastic disutility.

3.3 Procurement-Cost-Normalized Welfare Return

The raw welfare derivative $dW_{ss}/d\Delta x_j$ alone is not a sufficient basis for cross-region policy comparisons: regions differ in how costly urban land supply actually is. In the Korean context, for instance, developable land near already-urbanized metropolitan areas is scarcer and more expensive to convert to urban use than land in less-developed regions, which our cost convention $c_j^x = \kappa q_{j,0}/(1-\beta)$ captures in reduced form (Section 2.3). A policy-relevant comparison therefore normalizes the welfare derivative by the per-unit procurement cost—utility is scaled by u_c and cost by $\kappa q_{j,0}/(1-\beta)$, so the ratio has a direct interpretation as consumption-equivalent welfare per dollar of procurement spending (a welfare return on investment, or ROI). Dividing (11) by $u_c \kappa q_{j,0}/(1-\beta)$ gives the *procurement-cost-normalized welfare return* as follows:

$$\Pi_j^W \equiv \frac{1-\beta}{u_c \kappa q_{j,0}} \cdot \frac{dW_{ss}}{d\Delta x_j} = \underbrace{\frac{1}{\kappa}}_{\substack{\text{direct land}\\ \text{channel}}} + \underbrace{\frac{1-\beta}{\kappa\beta} \cdot \frac{1}{q_{j,0}} \cdot \frac{dK_{ss}}{d\Delta x_j}}_{\substack{\text{capital accumulation}\\ \text{channel}}} + \underbrace{\frac{\lambda}{\kappa(1+\lambda)} \cdot \frac{1}{q_{j,0}} \cdot \frac{dY_{ss}}{d\Delta x_j}}_{\text{externality channel}}. \quad (14)$$

The direct land channel equals $1/\kappa$ identically across regions, by the common-κ assumption of Section 2.3. The ROI interpretation is therefore useful: Π_j^W measures the gross consumption-equivalent welfare gain generated by one dollar of procurement spending in region j. In what follows, we use Π_j^W to rank regions by their welfare return per procurement dollar. Because the direct land channel is region-invariant, cross-region differences in Π_j^W are governed entirely by the capital accumulation and externality channels,

$$\Pi_j^W - \Pi_\ell^W = \frac{1-\beta}{\kappa\beta}\left[\frac{1}{q_{j,0}}\frac{dK_{ss}}{d\Delta x_j} - \frac{1}{q_{\ell,0}}\frac{dK_{ss}}{d\Delta x_\ell}\right] + \frac{\lambda}{\kappa(1+\lambda)}\left[\frac{1}{q_{j,0}}\frac{dY_{ss}}{d\Delta x_j} - \frac{1}{q_{\ell,0}}\frac{dY_{ss}}{d\Delta x_\ell}\right]. \quad (15)$$

Since both Π_j^W and Π_ℓ^W are consumption-equivalent welfare ROIs, their difference directly measures how many additional consumption-equivalent dollars of welfare each procurement dollar generates in the long run when allocated to region j instead of region ℓ.

This is the central theoretical benchmark of the paper. At the margin of the pre-policy steady state, the direct land channel return per dollar of procurement cost is equalized across regions—a consequence of the direct-price term combined with our common-κ cost convention, in which the per-unit procurement cost is proportional to land's pre-policy asset value with a region-invariant factor κ. Together these imply "location does not matter" for infinitesimal land supply at the pre-policy steady state. In our setting, however, the welfare-relevant object is evaluated at the *new* long-run steady state, where aggregate capital and regional allocations have fully re-adjusted to the land supply. This pushes the welfare return away from the uniform direct-channel benchmark: the steady state itself shifts, and the cross-region ranking is governed by two general-equilibrium forces—how much aggregate capital each unit of land in region j induces, and how much additional output it generates beyond what is priced by markets. The externality channel vanishes at $\lambda = 0$, reducing to the pure capital accumulation channel; at $\lambda > 0$, it reinforces allocation toward regions where the output response is large.

3.4 Optimal Allocation: Linearized Problem

With the intuition developed in the preceding subsections—a regional land supply triggers a chain of cross-region general-equilibrium adjustments summarized by the aggregate-capital response $dK_{ss}/d\Delta x_j$, which in turn determines the procurement-cost-normalized welfare return Π_j^W—we can characterize the optimal allocation analytically in the linearized limit. Replacing $W_{ss}(x_0 + \Delta x)$ with its first-order Taylor expansion around $\Delta x = 0$ converts the exact nonlinear government problem (10) into a linear program in $\{\Delta x_j\}$:

$$\max_{\Delta x \geq 0} \quad \sum_j \frac{\partial W_{ss}}{\partial \Delta x_j}\bigg|_{\Delta x = 0} \Delta x_j \quad \text{s.t.} \quad \sum_j \frac{\kappa\, q_{j,0}}{1-\beta} \Delta x_j \leq B, \tag{16}$$

whose optimum admits the following clean characterization.

Theorem 3 (Optimal allocation—linearized problem)**.** *An allocation $\{\Delta x_j^*\}$ solves (16) if and only if there exist a budget multiplier $\tilde{\mu} \geq 0$ and non-negativity multipliers $\{\tilde{\eta}_j \geq 0\}$ such that*

$$\Pi_j^W - \tilde{\mu} + \tilde{\eta}_j = 0, \qquad \tilde{\eta}_j\, \Delta x_j^* = 0, \qquad \tilde{\mu}\left(B - \sum_j \frac{\kappa\, q_{j,0}}{1-\beta} \Delta x_j^*\right) = 0, \qquad \forall j, \tag{17}$$

where Π_j^W is the pre-policy procurement-cost-normalized welfare return defined in (14). For any two active regions j, ℓ with $\Delta x_j^*, \Delta x_\ell^* > 0$, $\Pi_j^W = \Pi_\ell^W = \tilde{\mu}$; and if a single region j^* strictly dominates all others in Π^W, the optimum of (16) is a corner with the entire budget allocated to j^*.

Proof. See Appendix B.3. □

The key takeaway is direct: Π_j^W, the welfare return per dollar of procurement cost, is the ranking criterion. Linearizing around the pre-policy steady state, the optimal allocation pours the entire budget into the region with the highest pre-policy Π_j^W (or splits the budget across regions tied at the maximum). Normalizing welfare gains for the per-unit procurement cost is essential to this ranking—unnormalized welfare gains would mechanically favor regions with expensive land regardless of how productive the underlying allocation is.

From the linearized characterization to quantitative policy analysis. Theorem 3 characterizes the optimum at the infinitesimal-supply limit $\Delta x \to 0$. In the real world, government-driven land supply policies provide massive amounts of land: the Second NTP, for example, delivered 610,433 housing units in nine newly built SMA cities at a cost exceeding 10% of 2004 GDP. At that scale, Π_j^W falls as more land is supplied to a region through diminishing returns, so given a real budget the welfare-optimal allocation may well spread across multiple regions rather than concentrate in a single one; the corner solution of the linearized problem then gives way to an interior multi-region allocation that equalizes ex-post Π^W across active regions.[7] The shape of this allocation is ultimately a quantitative question, which we address via the counterfactual analysis using the structural model estimated from Korean regional data in Section 6.2.

4 South Korea: Institutional Details

4.1 Administrative Geography

South Korea is divided into 17 top-level administrative regions, shown in Figure 1a. We refer to them collectively as "provinces."[8] At the center is the Seoul Metropolitan Area (SMA), comprising Seoul and its two neighbors, Incheon and Gyeonggi. We drop two provinces, Sejong and Jeju, from the analysis: Sejong is an administrative capital hosting the central government, established in 2012, and Jeju is an island off the southern coast. Together they hold about 2% of the national population, and their specialized roles and distinct geography place them outside the scope of our analysis. We then assign the remaining major urban regions, such as Busan and Daegu, to the "Metro" group, and all other provinces to the "Rural" group. The resulting 15

[7] The exact KKT conditions for the nonlinear (10) evaluate $\partial W_{ss}/\partial \Delta x_j$ at the optimum Δx^*; because post-policy rents $\{q_j^*\}$ generally differ across regions and from $\{q_{j,0}\}$, the region-invariant direct land channel of (14) does not carry over exactly to the ex-post optimum.

[8] The 17 comprise eight Provinces, seven Metropolitan Cities, and two Special Self-Governing Provinces. These designations carry some administrative differences, but they play little role in our analysis.

provinces constitute our sample. Table 1 reports the population distribution across these groups.

Table 1: Population Distribution (2021)

Group	Region Name	Population Millions	%
Seoul Metropolitan Area (SMA)	Seoul, Incheon, Gyeonggi	26.0	(50.4)
	Seoul	9.5	(18.4)
	Incheon	2.9	(5.7)
	Gyeonggi	13.6	(26.3)
Metro	Busan, Daegu, Gwangju, Daejeon, Ulsan	9.8	(18.9)
Rural	All Provinces except Gyeonggi	14.8	(28.7)
Excluded	Sejong, Jeju	1.0	(2.0)

Notes: Population is measured in millions; shares are percentages of the national population.

The SMA is our primary region of interest. Seoul has been the national capital for over six centuries and the core of the country's economic development; Incheon and Gyeonggi urbanized more recently, drawn into Seoul's orbit as the city expanded after the Korean War. The Metro centers have historically served as economic hubs for their surrounding regions, with high population densities and the most developed urban amenities in their respective areas.

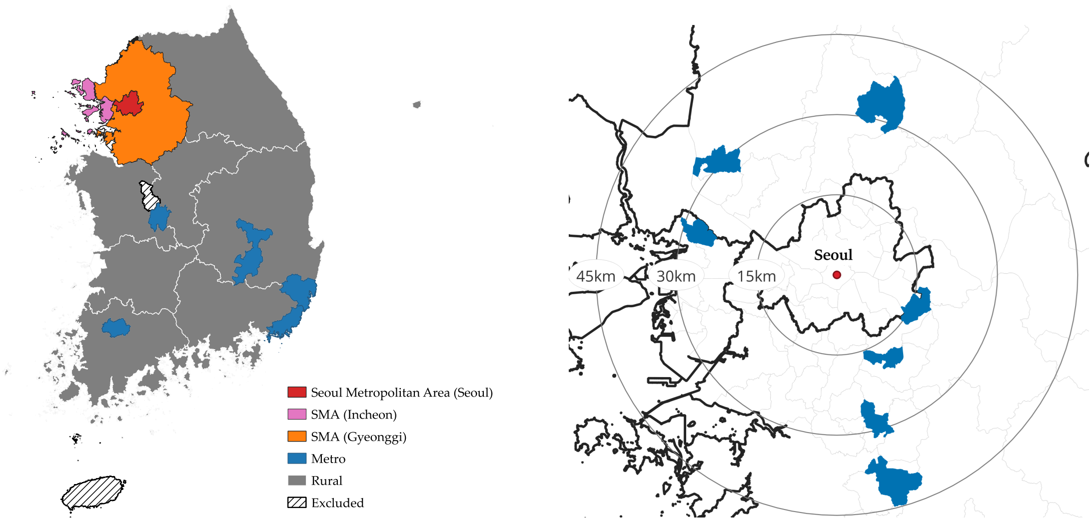

(a) Administrative Divisions (b) The Second New Town Projects

Figure 1: Administrative Divisions and New Town Projects

4.2 Concentration in the Seoul Metropolitan Area

South Korea's rapid growth since the 1970s has produced a striking concentration of population and economic activity in the Seoul Metropolitan Area. Between the

1970s and early 2000s, propelled by a combination of industrial policies, extensive infrastructure investments, and agglomeration, the SMA absorbed a steady influx of migrants who were disproportionately young and skilled, while many regions in the Rural group experienced population decline as well as demographic aging. Even the Metro group, including major cities such as Busan and Daegu, was no exception, underscoring the SMA's unique gravitational pull. (Table 1)

More than half of the nation's residents now live in this region, a share that continues to rise (Figure 2a). The SMA features a large labor market, dense supplier networks, and close proximity to both public and private headquarters. These synergies reinforce a virtuous cycle of concentrated growth. Gyeonggi, once predominantly rural, has evolved into a manufacturing and R&D powerhouse, home to semiconductor giants such as Samsung and SK Hynix. Its per capita GRDP (Gross Regional Domestic Product) relative to Seoul has remained broadly stable or even increased over the past two decades, while the Metro and Rural groups have experienced sharp relative declines (Figure 2b). Even Ulsan, the Metro city with the highest GRDP per capita and home to Hyundai Motors, shared in this relative decline.

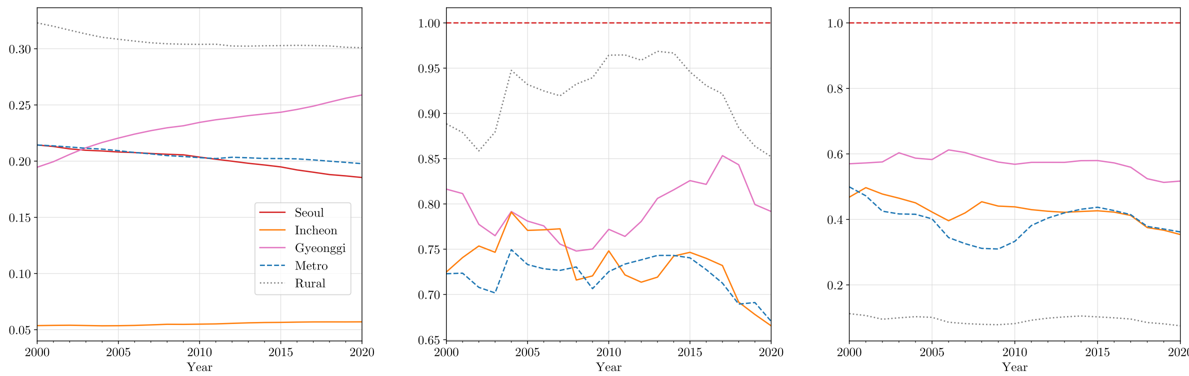

(a) Population share

(b) GRDP per capita

(c) Housing price

Figure 2: Regional Economic Trends

Notes: GRDP per capita and housing prices are normalized by Seoul.

Housing-price trends further illustrate the SMA's dominance. Between 2000 and 2020, housing prices in most regions fell relative to those in Seoul (Figure 2c). Gyeonggi is a notable exception: its relative housing prices remained broadly stable, reflecting its role as a feasible alternative to Seoul. This divergence points not only to the relative decline of non-SMA regions but also to spatial misallocation. High and rising housing costs in the SMA make it increasingly difficult for workers to relocate there, even as employment opportunities and infrastructure continue to concentrate in the region. We observe such spatial disparities in other developed countries as well, but

South Korea stands out.[9] South Korean policymakers responded with large-scale land supply initiatives.

4.3 New Town Projects

The South Korean government's New Town Projects serve as a key policy response to soaring housing prices and congestion in the SMA. Initiated in the late 1980s, these projects aimed to supply "fully serviced" urban land, complete with roads, utilities, and public amenities, on what were previously rural or undeveloped tracts near Seoul. The First New Town Project, designated in 1989, built five planned cities within roughly 20km of Seoul, ultimately providing 292,000 housing units. The Second New Town Project, designated from the early 2000s and implemented over the following two decades, further expanded the radius to 40 kilometers (Figure 1b) and delivered an additional 610,433 units in nine new cities at a cost of 95 trillion KRW (around 64 billion USD), exceeding 10% of South Korea's 2004 GDP.[10] The Third New Town Project, announced in 2018 and currently in progress, envisions 178,000 more units within approximately 20 kilometers of the capital.

Table 2: The 2nd New Town Project: Land and Housing supplied and Expenditure

Region	Government spending (tril. KRW)	Housing units	Land (km^2)			
			Total	Residential	Commercial	Industrial
SMA	94.9	610,433	49.6	36.6	5.8	7.2

Notes: Authors' calculation based on the official documents.

Unlike private-sector housing developments, which are typically carried out as scattered construction projects, the New Town Projects encompass entire "city packages." Each development integrates not only residential blocks but also transportation infrastructure, utilities, and commercial districts, effectively transforming large swaths of farmland or unused parcels into new viable urban land. This approach is facilitated by the Residential Site Development Promotion Act, a legal instrument that

[9]For example, the United Kingdom is known for spatial disparities and recently introduced a comprehensive plan, "Levelling Up", to address it. Greater London's population share is less than 15% as of 2022 (Office for National Statistics, 2022), which is significantly lower than the SMA's share of South Korea's population.

[10]The nine cities lie within the SMA, listed here with their host cities in parentheses: Pangyo (Seongnam), Dongtan (Hwaseong), Hangang (Gimpo), Unjeong (Paju), Gwanggyo (Suwon), Yangju, Wirye (Seongnam), Godeok (Pyeongtaek), and Geomdan (Incheon). Two of them, Dongtan and Yangju, each comprise two development districts (Dongtan 1 and 2; Okjeong and Hoecheon), so the official designation counts eleven SMA districts. Including two further districts outside the SMA brings the nationwide total to twelve districts and 666,000 units.

grants planning agencies the right of eminent domain upon designating a development site. Approval of a development plan simultaneously legalizes the corresponding rezoning, thus bypassing standard regulatory prerequisites. Empowered by the Act, the NTPs circumvent two critical obstacles—fragmented land ownership and the restrictions imposed by Restricted Development Zones (RDZs) or greenbelt designations—that make incremental upzoning nearly impossible. Moreover, although the government publicly disclosed its overarching plan to build new cities from the outset, the specific site locations remained confidential until the official announcement, creating a shock that was persistent in effect yet unanticipated for the designated regions.

This study focuses on the Second New Town Project, the largest of Korea's three New Town initiatives. The project provides a useful setting for studying optimal land supply in a regulated housing market. Unlike land-use deregulation, it created serviced urban land near the Seoul Metropolitan Area through public land assembly, rezoning, and infrastructure provision, but only by incurring substantial procurement and development costs. While anecdotal accounts credit the New Towns with alleviating housing shortages and moderating housing-price growth in the SMA, there has been little formal assessment of their long-run welfare effects. We provide the first quantitative evaluation of the Second NTP and ask whether its actual allocation maximized welfare per unit of public expenditure relative to optimal counterfactual allocations using the same budget.

5 Quantitative Analysis

5.1 Data

We assemble annual province-level data for 15 Korean provinces over 2004–2019. Population and employment come from the Economically Active Population Survey of Statistics Korea. We use the population aged 15 and older as the national population measure for normalization, and we use regional employment headcounts from the same survey for n_{jt}. Regional output, y_{jt}, is real gross regional domestic product from Statistics Korea's Regional Income Survey. For house prices, p_{jt}, we anchor each province's all-type average nominal house price to its 2015 level from the Korea Real Estate Board, chain the regional nominal house-price index over 2004–2019, and deflate by the consumer price index to obtain real house prices. These variables are constructed to be comparable to the objects used in HOP.

The land input, x_{jt}, is based on Cadastral Statistics, which reports detailed land areas by statutory land type. This differs from HOP, which measures state urban land by multiplying total available land acreage by an urban-land share from the U.S. Census. Our definition of urban land is based on the administrative categories in Korean land-use data. Specifically, we define urban land as the sum of "Land for Building and Commercial Use" and "Land for Industrial Use." This definition corresponds to the model's productive urban land: land that can be used for housing, commercial activity, and industrial production. We apply the same definition when mapping the 2nd New Town Project into the counterfactual simulations, so the policy land increment is measured consistently with the baseline land stock.

5.2 Calibration

We use Korean parameter values when they are available and otherwise follow the literature. We set $\beta = 0.9603$ to match the average 10-Year Treasury yield between 2004 and 2019. We set the depreciation rate to $\delta = 0.03$, a value commonly used in Korean general equilibrium models, including Bae (2013) and Song (2014). For the Frisch labor supply elasticity, we set $\gamma = 0.96$ based on the estimates of Moon and Song (2016).

For final-goods production, we follow HOP and set the land income share to 0.05, using the estimate in Valentinyi and Herrendorf (2008), because comparable Korean estimates are not available. We set the labor share to $\chi = 0.647$ and the capital share to $\theta = 0.303$, matching the average Korean labor income share over 2004–2019 while leaving the remaining 0.05 share for land. For housing production, we set the land input share to the scalar value $1 - \xi = 0.35$ in every region, following empirical work using South Korean data (Lee, 2024; Kim, 2020). This is consistent with the French-data estimate in Combes et al. (2021). For the procurement cost scalar κ, we use 2.09, taken from the Ministry of Land, Infrastructure, and Transportation (MOLIT) administrative dataset. Lastly, we use 0.03 for the production externality λ following HOP.

5.3 Model-inferred Regional Fundamentals

We follow Herkenhoff et al. (2018) and use the model to recover three regional fundamentals: total factor productivity A, amenities a, and land-use restriction parameter α. Because our main objective is to evaluate costly land-supply policies rather than to develop a new measurement framework for regional fundamentals, we keep this subsection brief.

Figure 3 reports the model-inferred values for the 2004 and 2019 long-run steady states. The inferred fundamentals are persistent across the two steady states: across

Table 3: Externally Calibrated Parameters

Name	Parameter	Value	Reference
Discount Factor	β	0.96	Average 10-Year Treasury yield (2004-2019)
Capital Depreciation Rate	δ	0.03	Bae (2013) and Song (2014)
Labor Supply Elasticity	γ	0.96	Moon and Song (2016)
Agglomeration Externality	λ	0.03	Herkenhoff et al. (2018)
Capital Share for Production	θ	0.303	Average Capital Income Share (2004-2019)
Labor Share for Production	χ	0.647	Average Labor Income Share (2004-2019)
Land Share for Production	$1-\theta-\chi$	0.05	Valentinyi and Herrendorf (2008)
Land Share for Housing	$1-\xi$	0.35	Lee (2024), Kim (2020)
Procurement cost factor	κ	2.09	MOLIT administrative reports

Notes: The table reports benchmark parameter values used in the quantitative model.

the 15 regions, the Spearman rank correlations are 0.89 for productivity, 0.92 for amenities, and 0.69 for α. These correlations indicate that, although levels change over time, regions' relative positions remain broadly stable. This persistence supports the use of the inferred fundamentals in the counterfactual exercises. The estimates of α are also economically intuitive. Seoul is the most restricted region, whereas the regions with the least restrictive inferred land-use regulations are all rural: Jeon-Nam, GyeongBuk, Gwangju, JeonBuk, and Kangwon. This pattern is consistent with Herkenhoff et al. (2018) and Hsieh and Moretti (2019), who find that land-use regulations are more restrictive in large, high-demand metropolitan areas in the United States. In the next subsection, we examine whether the inferred fundamentals are aligned with empirical proxies.

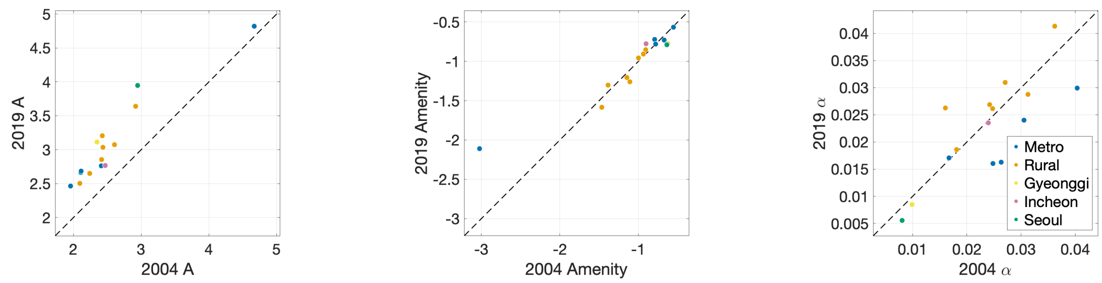

Figure 3: Model-inferred Regional Fundamentals (2004-2019)

5.4 Model Validation

The model matches the regional data on house prices, population, and labor productivity, following HOP, by construction. This leaves the remaining moments to validate the model against untargeted moments. We first compare model-implied land prices to observed land prices, because land prices are the key cost object in the government problem. We then compare the model-inferred values of amenities and land-use restrictions to empirical proxies as diagnostics.[11] We compare our estimates with quality-of-life measures and actual land-use restriction levels in South Korea by region.

To that end, we construct a separate land-price series. The South Korean government collects and discloses transaction-level land records since 2006. We collect all land transactions by province in 2010 and compute transaction prices per m^2 for land where residential, commercial, or factory construction is legally permitted.[12] We then compute province-level average land prices in 2010 and use regional nominal land-price indices from the Korea Real Estate Board to extend the series over time. Finally, we deflate nominal land prices by the consumer price index. This series is not used as a calibration target; it is used below as an untargeted validation moment for the model-implied land prices.

Figure 4 compares model-predicted land prices q_j with observed land prices in 2004 and 2019. For each year, both model and data land prices are normalized by Seoul. The model matches the data closely, both in level and fit (slope$=1.02$, $R^2 = 0.88$). Since land prices are not targeted, the comparison provides an out-of-sample validation of whether the model generates reasonable regional land values. The largest underpredictions are for Gyeonggi, the country's largest province, whose internal heterogeneity is substantial: its northern areas are close to the Demilitarized Zone and relatively less developed, while its southern areas include Suwon and major semiconductor clusters associated with Samsung and SK Hynix.[13]

Table 4 reports that the model-inferred amenities and land-use regulations are positively correlated with the empirical proxies.[14] As a proxy for amenities, we take the subjective satisfaction and quality of life indexes (2019) from the National Balanced-Development Information System (NABIS). The index has 22 subjects, such as "satisfaction with living in the region", "satisfactory natural environment", "public safety", "high-quality primary and secondary education." Our QoL measure is the first principal component of all subjects. The correlation is positive and statistically significant, for both Pearson and Spearman correlations.

[11]HOP compares the model-inferred TFP, amenities, and land-use restrictions. There is no widely accepted regional productivity estimate for South Korea, so we report the correlations for amenities and land-use restrictions.

[12]We exclude mountain land and land on which construction is legally prohibited for environmental conservation.

[13]With more spatially disaggregated data, these within-province differences could be modeled more directly.

[14]To the best of our knowledge, there are no widely accepted province-level productivity estimates in South Korea.

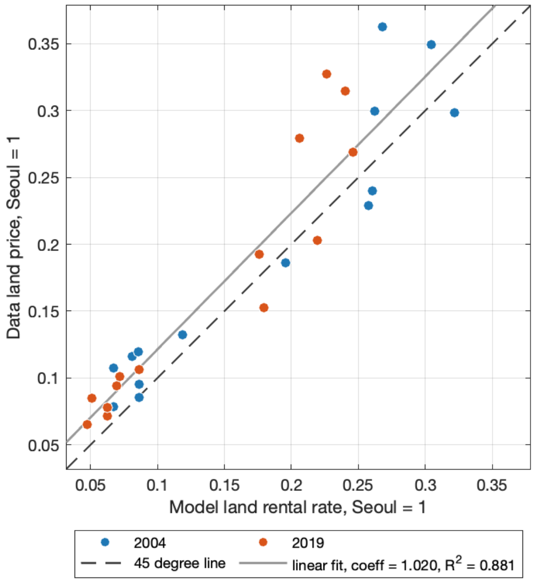

Figure 4: Model Validation: Observed (untargeted) versus Model-Predicted Land Prices

For land-use regulation, our parameter α_j is the fraction of urban land that may legally be utilized. A direct empirical counterpart is difficult to construct, because α_j is an aggregate share with no single regulatory instrument behind it. Still, popular building codes such as Floor Area Ratio (FAR) and Building Coverage Ratio (BCR) are in the same spirit because they are defined as ratios with respect to the land area. We obtained area-weighted averages of the ratios at the regional level during 2012-2019.[15] The model and data measures share the interpretation that any higher values imply more lenient restrictions. We normalize both the model and data land-use restrictions to Seoul for ease of comparison.

The model-inferred α_j is positively correlated with the observed land-use ratios across both measures. The Pearson correlations between the model-inferred α_j and observed land-use ratios are 0.600 for BCR, 0.530 for FAR, and 0.621 for the average of the two measures; the correlations are significant at the 5% level except for the FAR's Spearman correlation. The table thus provides supportive evidence that the inferred fundamentals are aligned with empirical proxies. These results used the 2019 estimates, because our empirical proxies are not available for 2004.

[15]We have permit-level data for FAR and BCR, but the time from permission to completion varies substantially across permits, so it is tricky to map a single year's statistic to our steady state. Since we are focusing on the steady states, we pooled the permits and legal regulations over time.

Table 4: Model Validation: Inferred Fundamentals versus Empirical Proxies

	Amenity	Land-use Regulation		
		Building Coverage Ratio	Floor Area Ratio	Average (BCR,FAR)
Pearson Correlation	0.505 (0.055)	0.600 (0.018)	0.530 (0.042)	0.621 (0.013)
Spearman Correlation	0.582 (0.023)	0.482 (0.069)	0.396 (0.143)	0.521 (0.046)

Notes: The amenity proxy is the first principal component of NABIS subjective satisfaction and quality-of-life indexes. The land-use proxy uses legal FAR and BCR for residential, commercial, and industrial urban land, area-weighted within category and region, then combined using a simple mean. Parentheses are p-values.

6 Counterfactual Experiments: Optimal Land Supply

6.1 Setup

The counterfactual exercises ask where new land should be supplied. We take the 2019 steady state, in which the actual 2nd New Town Project (NTP) land supply has been realized in the data, and proceed in two steps. First, we construct a no-NTP economy by removing all policy-supplied land from the Seoul Metropolitan Area (SMA), which isolates the long-run effect of the actual NTP. Second, with the no-NTP economy as the common baseline, we compare the actual NTP allocation against the ex ante optimal one. We define the ex ante optimal allocation as the one that ranks regions by their 2004 procurement-normalized welfare return, Π^W, and assigns the entire NTP development budget to the highest-return region. The allocation is ex ante because the ranking uses only the 2004 economy, the information available when the 2nd NTP was announced and implementation began.

6.2 New Town Project: Counterfactual Scenarios

All results in this section are reported relative to the 2019 no-NTP economy. For the actual NTP, this comparison measures the long-run effect of the implemented policy, supplying land to Seoul, Gyeonggi, and Incheon. For the ex ante optimal allocation, it measures the effect of reallocating the same development budget according to the 2004 procurement-normalized welfare returns; in the benchmark case, this rule assigns the full NTP budget to Daejeon. The reported gains for the ex ante optimal NTP therefore measure the long-run effect of allocating the same development budget to Daejeon in the 2019 no-NTP economy.

Table 5: Welfare Effects of Actual and Ex Ante Optimal NTP, Relative to No NTP

Scenario	Welfare CE (%)	GDP (%)	Capital (%)
Actual NTP	0.280	0.226	0.191
Ex ante optimal NTP	0.350	0.353	0.379

Notes: All changes are reported relative to the no-NTP counterfactual. The table evaluates the actual and ex ante optimal finite land allocations in the full nonlinear steady-state model. Welfare is a consumption-equivalent change; GDP and capital are percentage changes in aggregate levels.

Aggregate Output and Welfare Table 5 reports the aggregate effects of the actual and ex ante optimal NTP, both relative to the no-NTP benchmark. The actual project raises welfare by 0.280% in consumption-equivalent terms. It also raises aggregate GDP by 0.226% and aggregate capital by 0.191%. The ex ante optimal allocation improves on the actual NTP when evaluated in the full nonlinear model. The welfare gain is 0.350%, compared with 0.280% under the actual project. The difference is 0.070 percentage points, so the actual NTP captures about 80% of the welfare gain available under the ex ante same-budget allocation. The gap is larger for output. The ex ante optimal NTP raises GDP by 0.353%, compared with 0.226% under the actual NTP, so the actual allocation captures about 64% of the ex ante optimal GDP gain. Aggregate capital rises by 0.379%, compared with 0.191% under the actual NTP; the actual allocation captures about 50% of the ex ante optimal capital gain.

Actual versus Ex Ante Optimal Allocation The actual 2nd NTP budget was concentrated in the SMA: 80.2% in Gyeonggi, 10.1% in Incheon, and 9.6% in Seoul. On the other hand, the 2004-informed allocation places the entire land supply in Daejeon, a non-SMA metropolitan city with the highest 2004 procurement-normalized welfare return Π^W. The actual SMA regions, Gyeonggi, Incheon, and Seoul, also ranked relatively high. This is a corner solution, not a "balanced-development" allocation that spreads the budget across many non-SMA regions. The result should be interpreted through the model's 2004 regional fundamentals and procurement costs. Daejeon, Daegu, and Seoul are close in procurement-normalized returns, but they remain ordered in the benchmark: $\Pi^W = 0.753$ for Daejeon, 0.746 for Daegu, and 0.742 for Seoul. Although the top procurement-normalized returns are numerically close, the same-budget optimum remains a corner solution. With the actual 2nd NTP budget, allocating all investment to Daejeon does not reduce its marginal return enough to make diversification optimal. In the benchmark, investment begins to spread across regions only when the budget is roughly 2.5 times larger than the actual budget. We will discuss the underlying mechanisms in the following subsection.

Table 6: Actual NTP versus Ex Ante Optimal NTP Allocation

Region	Actual budget share (%)	Ex ante opt. budget share (%)	Actual land increment	Ex ante opt. land increment	2004 Π^W rank
Daejeon	–	100.0	–	0.0318	1
Gyeonggi	80.2	–	0.0244	–	7
Incheon	10.1	–	0.0030	–	5
Seoul	9.6	–	0.0007	–	3

Notes: Budget shares sum to the total 2nd NTP land-development budget. The ex ante optimal allocation is selected using 2004 procurement-normalized returns and evaluated in the 2019 no-NTP economy. The land increments are measured in acres per capita.

6.3 Unpacking the NTP Scenarios: Mechanisms and Regional Outcomes

We interpret the aggregate results in Table 5 through the lens of the theory. Using the decomposition equations such as (12), (13), and (14), we show the underlying mechanisms of the ex ante optimal allocation. Thus the decomposition is a ranking and mechanism exercise, while the aggregate table is the full-model evaluation of the finite counterfactuals.

Table 7: 2004 Ex Ante Procurement-Normalized Welfare Return Decomposition

Policy allocation	Π^W	Direct	Capital accumulation	Externality
Actual NTP	0.665	0.483	0.163	0.019
Ex ante optimal NTP	0.753	0.483	0.242	0.028

Notes: Rows are procurement-budget-weighted averages of the 2004 regional components of Equation (14). Entries are procurement-normalized welfare-return units, not consumption-equivalent percentages in Table 5. The actual row uses model procurement-budget shares for the actual NTP allocation; the ex ante optimal row places all budget in Daejeon.

Welfare Decomposition Table 7 decomposes the ex ante procurement-normalized return. The actual NTP averages to $\Pi^W = 0.665$, while the ex ante optimal NTP reaches 0.753. The direct term is 0.483 in both rows because the direct land return is normalized by a procurement cost proportional to the same pre-policy land rent. The gap therefore comes from general-equilibrium adjustment: the capital accumulation component rises from 0.163 to 0.242, and the externality component rises from 0.019 to 0.028. The production externality matters, but the main reason Daejeon ranks first is the larger capital accumulation return per procurement dollar.

Table 8: 2004 Ex Ante Capital Accumulation Decomposition

Policy allocation	Capital accumul.	Prod. land-capital comp.	Housing land-capital subst.	Fixed-employment reallocation	Aggregate labor supply
Actual NTP	0.163	0.060	-0.148	0.011	0.241
Ex ante optimal NTP	0.242	0.064	-0.139	-0.139	0.456

Notes: Entries are procurement-budget-weighted 2004 capital accumulation components in procurement-normalized welfare-return units. The components are calculated by applying the capital-response decomposition in Equation (12) and the labor split in Equation (13) to the capital accumulation term in Equation (14). The components sum to the capital accumulation column up to rounding.

Table 8 shows where the capital accumulation gap comes from. Production land-capital complementarity and housing land-capital substitution are similar across the two allocations: 0.060 versus 0.064 on the production side, and -0.148 versus -0.139 on the housing side. The difference is the labor margin. The ex ante optimal allocation has a much larger aggregate labor-supply component, 0.456 rather than 0.241, even though its fixed-employment reallocation term is negative at -0.139. Supplying land to Daejeon reallocates workers away from some more capital-intensive regions at fixed national employment, but it generates a stronger aggregate employment response.

The decomposition in Theorem 2, together with equation (13), isolates the two channels behind this result. Holding aggregate employment fixed, supplying land to Daejeon reallocates workers away from regions where labor is paired with more capital. Because the rental rate of capital is common across regions, this reallocation lowers the economy-wide demand for capital and reduces the fixed-employment component. The key offsetting force is the extensive margin. Additional land in Daejeon relaxes the local housing constraint and raises the net return to working there. More workers can therefore move to Daejeon at lower housing cost, while the regions they leave behind become less congested. This adjustment lowers the economy-wide cost of accommodating workers and increases aggregate labor supply. By contrast, an SMA-centered expansion directs land toward already dense and expensive locations, where the same procurement budget translates into a smaller increase in serviced urban land, leading to a weaker aggregate labor-supply response. This extensive-margin response is the main capital accumulation channel through which the optimal allocation dominates the actual SMA-centered allocation.

Table 9: 2004 Π^W Decomposition (Equation 28): Top Five Regions and Seoul

Region	Π^W rank	Π^W	$dW/d\Delta x$	dW/dx rank	q	Direct	Capital Accumul.	Externality.
Daejeon	1	0.753	3.327	5	0.204	0.483	0.242	0.028
(Seoul=1)		(1.02)	(0.26)		(0.26)			
Daegu	2	0.746	3.352	4	0.208	0.483	0.238	0.024
(Seoul=1)		(1.01)	(0.26)		(0.26)			
Seoul	3	0.742	12.718	1	0.793	0.483	0.241	0.018
(Seoul=1)		(1.00)	(1.00)		(1.00)			
Busan	4	0.700	3.867	2	0.255	0.483	0.193	0.024
(Seoul=1)		(0.94)	(0.30)		(0.32)			
Incheon	5	0.696	3.631	3	0.241	0.483	0.184	0.028
(Seoul=1)		(0.94)	(0.29)		(0.30)			

Notes: The table implements the procurement-normalized welfare decomposition in Equation (14). It is evaluated in the 2004 design environment and sorted by 2004 Π^W rank. The columns for $dW/d\Delta x$ and q show the raw marginal welfare return and the flow land price that enter the procurement-normalized return. Rank columns are descending ranks across all 15 regions. Parentheses report ratios with respect to Seoul, normalized to one, for Π^W, $dW/d\Delta x$, and q. The summary keeps the top five regions plus Seoul.

Decision Rule Table 9 explains why the ex ante allocation selects Daejeon through the lens of the model. The table reports the 2004 procurement-normalized welfare return, Π^W, using the decomposition in Equation (14). The direct-return term is identical across regions, following the envelope logic in the static equilibrium. The ranking is therefore driven by the capital accumulation term, the externality term, and procurement cost.

Daejeon has the highest procurement-normalized regional investment return, Π^W, 0.753, which is about 2% higher than Seoul. It combines a raw marginal welfare return of $dW/d\Delta x = 3.327$ with a moderate land price, $q = 0.204$, and a high capital accumulation contribution. Daegu is very close behind, with $\Pi^W = 0.746$. Seoul illustrates the role of land price and normalization. It has the largest raw marginal welfare return among the regions shown, $dW/d\Delta x = 12.718$, but its land price is also much higher, $q = 0.793$, so procurement-cost normalization pushes it below Daejeon and Daegu. Seoul nevertheless remains third because its capital accumulation component, 0.241, is nearly as large as Daejeon's under the scalar benchmark. The comparison among Daejeon, Daegu, and Seoul shows that the ex ante land-supply rule is not simply to build in the most productive or most congested region; it depends on the full procurement-normalized general-equilibrium return.

The top five regions are all urban or near-urban alternatives: Daejeon, Daegu, Seoul, Busan, and Incheon. In this benchmark, the procurement-normalized return rewards places that combine moderate procurement costs with strong capital accumulation gains. Daejeon wins because it has the highest such combined return, not because it has the highest raw marginal value of land. Appendix Table A.1 reports Π^W, $dW/d\Delta x$, and q for all 15 provinces.

Regional Fundamentals and Π^W. Table 10 summarizes how the procurement-cost-normalized welfare return, Π_j^W, is related to regional fundamentals. The table reports both unconditional rank correlations and multivariate regressions that include regional productivity, amenities, and the effective urban land stock jointly. These results describe equilibrium relationships within the calibrated economy, not causal effects of the individual fundamentals.

Regional productivity, A_j, has two opposing implications for the welfare return to land supply. More productive regions generate larger raw marginal welfare gains from additional urban land, as reflected in the strong positive conditional relationship between A_j and $dW^{ss}/d\Delta x_j$. At the same time, productivity also raises the equilibrium land rent $q_{j,0}$ and therefore the cost of acquiring and developing land. Once the welfare gain is normalized by this procurement-cost anchor, productivity is no longer a strong positive predictor of Π_j^W: its unconditional rank correlation with Π_j^W is negative, while its conditional coefficient is small. Thus, high productivity raises the gross return to additional land, but much of this advantage is offset by the higher cost of supplying land in productive regions.

Amenities, a_j, exhibit a different pattern. Higher amenities are positively associated with the land rent and the raw welfare gain, but they also remain strongly positively associated with Π_j^W after procurement costs are taken into account. This pattern suggests that the welfare benefits associated with amenities are not fully capitalized into the local cost of land. In the model, a high-amenity region is more attractive to workers, so supplying additional urban land there can induce a larger aggregate labor-supply response and, through labor–capital complementarity, a larger long-run expansion in capital, output, and welfare. This interpretation is consistent with the decomposition in Table 9, in which differences in labor adjustment are an important source of variation in the capital accumulation return.

The effective urban land stock, $\alpha_j x_j$, is negatively associated with Π_j^W. This relationship, however, should not be interpreted as showing that regions with less effective land necessarily generate larger raw welfare gains or face lower land prices: in the multivariate regressions, $\alpha_j x_j$ is not strongly related to either $dW^{ss}/d\Delta x_j$ or $q_{j,0}$ separately. Rather, the result indicates that the welfare gain relative to the local procurement-cost anchor tends to be higher in regions with a smaller existing stock of effective urban land. Because this ratio reflects the full general-equilibrium response, the relationship cannot generally be attributed to a single adjustment channel.

Taken together, these findings qualify the conventional prescription that new urban land should necessarily be supplied to the most productive region. That prescription focuses on the gross benefit of relaxing a spatial constraint but does not account for the high acquisition and development costs that are often capitalized into land values in productive locations. For costly active land-supply policies, the relevant criterion is instead the general-equilibrium welfare gain per dollar of procurement cost. Regions with strong amenities may therefore offer more cost-effective locations for new urban development than regions with the highest productivity alone, because they can generate a stronger labor-supply and long-run adjustment response relative to the cost of the land.

Table 10: Associations between 2004 Π^W and Regional Fundamentals

	Policy-return measure		
	Π^W	q	dW/dx
Panel A. Spearman rank correlations			
A	-0.461	0.036	-0.086
Amenity	0.875	0.650	0.725
α	-0.332	-0.396	-0.371
x	-0.461	-0.354	-0.461
αx	-0.625	-0.686	-0.750
Panel B. OLS regressions			
A	0.032	0.609***	9.704***
	(0.026)	(0.081)	(1.329)
Amenity	0.160***	0.659***	11.022***
	(0.028)	(0.086)	(1.419)
$1000 \times \alpha x$	-0.027***	0.000	0.004
	(0.004)	(0.012)	(0.197)
Constant	0.789***	-0.667***	-10.339***
	(0.048)	(0.147)	(2.410)
N	15	15	15
R^2	0.958	0.883	0.883

Notes: Panel A reports Spearman rank correlations across the 15 regions in the 2004 steady state. $\alpha_j x_j$ is the effective urban land stock. Panel B reports OLS coefficients from regressions of each policy-return measure on the listed fundamentals; robust standard errors in parentheses. *** denotes significance at the 1 percent level.

Employment Shares, Housing Prices and Regional Output Figure 5 summarizes the regional outcomes under the actual and ex ante optimal allocations. All panels report Actual NTP and Ex Ante Optimal NTP relative to the No NTP benchmark. The actual NTP raises employment shares in Gyeonggi and Incheon, shifting activity toward the SMA provinces that received most of the project land. The ex ante optimal NTP instead shifts the growth margin to the Metro group because Daejeon receives the full budget. Metro employment rises by 2.03 percentage points relative to the No

NTP benchmark, and Metro GRDP per capita rises by 10.7%. Both allocations lower SMA housing prices relative to the No NTP benchmark. The main regional implication is therefore not that the actual NTP failed: it generated positive aggregate gains and lowered SMA housing prices. Rather, the ex ante optimal allocation generates larger aggregate welfare and output gains while counteracting SMA concentration by moving the policy-induced growth margin outside the SMA.

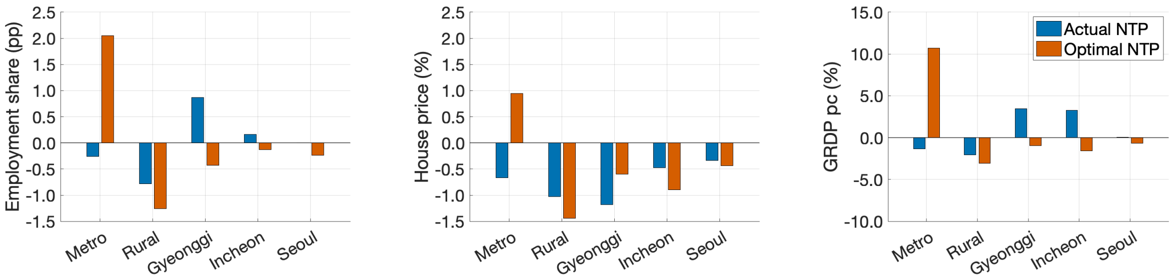

(a) Employment shares (b) Housing prices (c) GRDP per capita

Figure 5: Regional Outcomes under Actual NTP and Ex Ante Optimal NTP, Relative to No NTP

Finally, the comparison should be interpreted as a steady-state counterfactual. The model includes production externalities, so the exercise is not mechanically missing agglomeration forces. But dynamic adjustment can still matter. If the same land had been supplied to Daejeon in 2004, migration, capital, local prices, and externalities could have evolved along a different transition path before reaching the 2019 evaluation environment. The steady-state comparison is therefore best read as evidence on the long-run allocation margin, not as a full dynamic history of the Korean regional economy.

7 Conclusion

This paper studies active urban land supply as a public-investment problem in a spatial economy. Regions where land is most constrained offer the highest marginal value from new supply, but they are also the most expensive places to supply it. High land prices signal high marginal land values, but they do not by themselves identify where public land investment delivers the largest welfare gain. The relevant criterion is the welfare return per dollar, which reflects both the marginal value of land and the general-equilibrium response through migration, housing prices, capital accumulation, and production externalities.

To our knowledge, this is the first paper to study active urban land supply as a budget-constrained spatial allocation problem and to solve for its optimal location in a quantitative general-equilibrium model. Korea's 2nd New Town Project provides a useful laboratory for this question. The actual SMA-centered project generated positive long-run aggregate gains: welfare, GDP, and capital are all higher than in the no-NTP counterfactual. At the same time, the optimal land-supply rule implies that the same budget could have produced larger steady-state gains outside the SMA. In the baseline counterfactual exercise, the welfare-maximizing allocation places the project in Daejeon and raises aggregate welfare and output more than the actual allocation, while redirecting population and capital accumulation away from further SMA concentration in the long run.

Our characterization of the optimal land supply rests on a stylized model. The quantitative result for the welfare-maximizing allocation may depend on model structure, but the key theoretical result does not: the marginal value of land alone cannot identify where active land supply should go. This is a standard principle of welfare economics, but one the land-supply literature has not applied to the location problem. The planner must weigh acquisition and development costs against the general equilibrium gains from new urban land, a trade-off inherent in any budget-constrained allocation problem.

Our analysis compares steady states after permanent land-supply changes and is silent on transition dynamics. The framework also abstracts from worker heterogeneity and other dimensions central to modern spatial economics and optimal spatial policy. Extending the analysis to study how optimal land supply interacts with transition dynamics, spatial redistribution policies, and fiscal instruments would be a natural avenue for future work.

References

Albouy, D., Ehrlich, G., 2018. Housing productivity and the social cost of land-use restrictions. Journal of Urban Economics 107, 101–120. doi:10.1016/j.jue.2018.06.002.

Allen, T., Arkolakis, C., 2014. Trade and the topography of the spatial economy. The Quarterly Journal of Economics 129, 1085–1140. doi:10.1093/qje/qju016.

Babalievsky, F., Herkenhoff, K.F., Ohanian, L.E., Prescott, E.C., 2026. The impact of commercial real estate regulations on U.S. output. URL: https://www.nber.org/papers/w31895, doi:10.3386/w31895. accepted at Review of Economic Studies. Also available as NBER Working Paper No. 31895.

Bae, B.H., 2013. The Role of Financial Factors in the Business Cycle and the Transmission of Monetary Policy in Korea. Working Paper 2013-30. Economic Research Institute, Bank of Korea. doi:10.2139/ssrn.2580595. in Korean.

Busso, M., Gregory, J., Kline, P., 2013. Assessing the incidence and efficiency of a prominent place based policy. American Economic Review 103, 897–947. doi:10.1257/aer.103.2.897.

Caliendo, L., Dvorkin, M., Parro, F., 2019. Trade and labor market dynamics: General equilibrium analysis of the China trade shock. Econometrica 87, 741–835. doi:10.3982/ECTA13758.

Combes, P.P., Duranton, G., Gobillon, L., 2021. The production function for housing: Evidence from France. Journal of Political Economy 129, 2766–2816. doi:10.1086/715076.

Department of Transport and Planning, Victoria, 2024. A 10-year plan for Melbourne's greenfields. State Government of Victoria. URL: https://www.planning.vic.gov.au/guides-and-resources/strategies-and-initiatives/housing/a-10-year-plan-for-melbournes-greenfields.

Desmet, K., Rossi-Hansberg, E., 2014. Spatial development. American Economic Review 104, 1211–1243. doi:10.1257/aer.104.4.1211.

Diamond, R., 2016. The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000. American Economic Review 106, 479–524. doi:10.1257/aer.20131706.

Donaldson, D., Hornbeck, R., 2016. Railroads and American economic growth: A "market access" approach. The Quarterly Journal of Economics 131, 799–858. doi:10

.1093/qje/qjw002.

Furman, J., 2015. Barriers to shared growth: The case of land use regulation and economic rents. Remarks at the Urban Institute. URL: https://obamawhitehouse.archives.gov/sites/default/files/page/files/20151120_barriers_shared_growth_land_use_regulation_and_economic_rents.pdf.

Ganong, P., Shoag, D., 2017. Why has regional income convergence in the U.S. declined? Journal of Urban Economics 102, 76–90. doi:10.1016/j.jue.2017.07.002.

Glaeser, E.L., 2014. Land use restrictions and other barriers to growth. Cato Online Forum. URL: https://www.cato.org/cato-online-forum/land-use-restrictions-other-barriers-growth.

Greaney, B., 2026. Housing constraints and spatial misallocation: Comment. American Economic Journal: Macroeconomics 18, 409–428. doi:10.1257/mac.20230141.

Herkenhoff, K.F., Ohanian, L.E., Prescott, E.C., 2018. Tarnishing the golden and empire states: Land-use restrictions and the U.S. economic slowdown. Journal of Monetary Economics 93, 89–109. doi:10.1016/j.jmoneco.2017.11.001.

Hsieh, C.T., Moretti, E., 2019. Housing constraints and spatial misallocation. American Economic Journal: Macroeconomics 11, 1–39. doi:10.1257/mac.20170388.

Kim, J.G., 2020. Apportioned prices of the land and single-residential house: A sales-price based approach. Journal of the Korean Cartographic Association 20, 57–69. doi:10.16879/jkca.2020.20.2.057.

Kline, P., Moretti, E., 2014. Local economic development, agglomeration economies, and the big push: 100 years of evidence from the Tennessee Valley Authority. The Quarterly Journal of Economics 129, 275–331. doi:10.1093/qje/qjt034.

Korsak, B.d., Pernelle, J., 2004. L'évaluation de la politique du logement dans les villes nouvelles. Rapport. Conseil général des Ponts et Chaussées, Ministère de l'Équipement, des Transports et du Logement. URL: https://documentation.insp.gouv.fr/insp/doc/SYRACUSE/123175/l-evaluation-de-la-politique-du-logement-dans-les-villes-nouvelles-rapport-etabli-par-bernard-de-kor. in French.

Koster, H.R.A., van Ommeren, J., 2019. Place-based policies and the housing market. The Review of Economics and Statistics 101, 400–414. doi:10.1162/rest_a_00779.

Lee, S.M., 2024. Empirical Study on the Relationship between Apartment Sale Prices and Land and Construction Costs. Master's thesis. Gachon University. Seongnam, Korea. URL: https://www.dbpia.co.kr/journal/detail?nodeId=T16946212.

master's thesis. Public metadata verified through DBpia listing.

Ministry of Housing, Communities and Local Government, 2024. Policy statement on new towns. GOV.UK. URL: https://www.gov.uk/government/publications/policy-statement-on-new-towns/policy-statement-on-new-towns.

Moon, W., Song, S., 2016. 노동공급 탄력성 추정. 노동경제논집 39, 35–51. In Korean.

Moretti, E., 2010. Local multipliers. American Economic Review 100, 373–377. doi:10.1257/aer.100.2.373.

Moretti, E., 2013. Real wage inequality. American Economic Journal: Applied Economics 5, 65–103. doi:10.1257/app.5.1.65.

Neumark, D., Simpson, H., 2015. Place-based policies, in: Duranton, G., Henderson, J.V., Strange, W.C. (Eds.), Handbook of Regional and Urban Economics. Elsevier. volume 5, pp. 1197–1287. doi:10.1016/B978-0-444-59531-7.00018-1.

Redding, S.J., Rossi-Hansberg, E., 2017. Quantitative spatial economics. Annual Review of Economics 9, 21–58. doi:10.1146/annurev-economics-063016-103713.

Roback, J., 1982. Wages, rents, and the quality of life. Journal of Political Economy 90, 1257–1278. doi:10.1086/261120.

Song, I.H., 2014. House price channel: Effects of house prices on macroeconomy. KDI Journal of Economic Policy 36, 171–205. doi:10.23895/kdijep.2014.36.4.171.

United Kingdom Parliament, 1946. New towns act 1946. Legislation.gov.uk. URL: https://www.legislation.gov.uk/ukpga/Geo6/9-10/68/enacted.

Urban Renaissance Agency, 2007. Tama New Town outline. Urban Renaissance Agency. URL: https://www.ur-net.go.jp/e/aboutus/profile/tama_new_town_en.html.

Valentinyi, Á., Herrendorf, B., 2008. Measuring factor income shares at the sectoral level. Review of Economic Dynamics 11, 820–835. doi:10.1016/j.red.2008.02.003.

A More Results

Table A.1: Welfare Objects by Region, Ranked by 2004 Π^W

Region	2004 Π^W	2004 Π^W rank	2004 dW/dx	2004 dW/dx rank	2004 q
Daejeon	0.753	1	3.327	5	0.204
Daegu	0.746	2	3.352	4	0.208
Seoul	0.742	3	12.718	1	0.793
Busan	0.700	4	3.867	2	0.255
Incheon	0.696	5	3.631	3	0.241
Gwangju	0.680	6	2.280	7	0.155
Gyeonggi	0.652	7	2.996	6	0.213
GyeongNam	0.619	8	1.261	9	0.094
Kangwon	0.617	9	0.916	10	0.069
JeonBuk	0.575	10	0.661	14	0.053
Chungbuk	0.575	11	0.797	13	0.064
Chungnam	0.546	12	0.800	12	0.068
GyeongBuk	0.542	13	0.805	11	0.069
JeonNam	0.465	14	0.538	15	0.053
Ulsan	0.423	15	1.890	8	0.207

Notes: The table reports the procurement-normalized welfare return Π^W, raw welfare return $dW/d\Delta x_j$, and model land price q_j evaluated in the 2004 steady state.

B Proofs

B.1 Proof of Theorem 1

Proof of Theorem 1. Since $x_{jt} = x_{j,0} + \Delta x_j$, $dx_j/d\Delta x_j = 1$ and hence $dW_{\text{ss}}/d\Delta x_j = dW_{\text{ss}}/dx_j$. Throughout this proof, all variables denote steady-state objects; for readability we drop the "ss" subscript and write W, K, Y for $W_{\text{ss}}, K_{\text{ss}}, Y_{\text{ss}}$. We derive the identity for dW/dx_j at a general steady state and specialize to the pre-policy steady state ($\Delta x = 0$, where $q_j = q_{j,0}$ and $c = c_0$) at the end.

Steady-state welfare is $W(x) = \frac{1}{1-\beta}\big[u(c(x)) + \sum_\ell \big(a_\ell\, n_\ell(x) - v(n_\ell(x))\big)\big]$, where u is a strictly concave, differentiable consumption utility, v is a convex, differentiable labor-disutility function, and $\{c(x), n_\ell(x)\}$ are competitive-equilibrium outcomes as functions of the land-endowment vector x (the baseline (1) uses $u(c) = \ln c$ and $v(n) = n^{1+1/\gamma}/(1+1/\gamma)$). We write $u_c \equiv u'(c)$ throughout; the derivation below uses only that u and v are differentiable with $u_c > 0$. Because W is not the household's value function but rather aggregate welfare evaluated at competitive-equilibrium outcomes, equilibrium prices themselves depend on x and the chain-rule terms in n_ℓ are not guaranteed to collapse via the envelope theorem. Chain-rule differentiation gives

$$\frac{dW}{dx_j} = \frac{1}{1-\beta}\left[u_c \frac{dc}{dx_j} + \sum_\ell (a_\ell - v'(n_\ell)) \frac{dn_\ell}{dx_j}\right]. \tag{18}$$

The household's first-order condition for labor supply to region ℓ at the competitive equilibrium is

$$a_\ell - v'(n_\ell) = u_c\,(p_\ell - w_\ell). \tag{19}$$

Substituting (19) into (18),

$$\frac{dW}{dx_j} = \frac{u_c}{1-\beta}\left[\frac{dc}{dx_j} + \sum_\ell (p_\ell - w_\ell) \frac{dn_\ell}{dx_j}\right]. \tag{20}$$

The steady-state resource constraint $c = Y - \delta K$, with $Y = \sum_j y_j$, gives $dc/dx_j = dY/dx_j - \delta\, dK/dx_j$.

We next derive an identity for dY/dx_j at the competitive equilibrium. With $y_\ell = A_\ell \bar{y}_\ell^{1+\lambda}$, the social partial derivatives are

$$\frac{\partial y_\ell}{\partial k_{y\ell}} = (1+\lambda)\,r, \quad \frac{\partial y_\ell}{\partial n_\ell} = (1+\lambda)\,w_\ell, \quad \frac{\partial y_\ell}{\partial x_{y\ell}} = (1+\lambda)\,q_\ell, \tag{21}$$

using firm FOCs $r = \theta y_\ell / k_{y\ell}$, $w_\ell = \chi y_\ell / n_\ell$, $q_\ell = \varphi y_\ell / x_{y\ell}$ (firm treats $\widehat{A}$ as given). Hence

$$\frac{dY}{dx_j} = (1+\lambda) \sum_\ell \left[r \frac{dk_{y\ell}}{dx_j} + w_\ell \frac{dn_\ell}{dx_j} + q_\ell \frac{dx_{y\ell}}{dx_j}\right]. \tag{22}$$

We simplify (22) using three identities.

(i) Aggregate capital. Differentiating $K = \sum_\ell (k_{y\ell} + k_{h\ell})$ with respect to x_j (and noting r is common across regions),

$$\sum_\ell r \frac{dk_{y\ell}}{dx_j} = r \frac{dK}{dx_j} - r \sum_\ell \frac{dk_{h\ell}}{dx_j}. \tag{23}$$

(ii) Regional land. Differentiating the region-ℓ land constraint $x_\ell = x_{y\ell} + x_{h\ell}$ and using $dx_\ell / dx_j = \delta_{\ell j}$ (Kronecker; $= 1$ if $\ell = j$, else 0) gives $dx_{y\ell}/dx_j = \delta_{\ell j} - dx_{h\ell}/dx_j$.

Weighting by q_ℓ and summing,

$$\sum_\ell q_\ell \frac{dx_{y\ell}}{dx_j} = q_j - \sum_\ell q_\ell \frac{dx_{h\ell}}{dx_j}. \tag{24}$$

(iii) Housing-firm identity. Differentiating the binding housing feasibility $n_\ell = k_{h\ell}^{\xi}(\alpha_\ell x_{h\ell})^{1-\xi}$ gives $dn_\ell = \xi(n_\ell/k_{h\ell})\, dk_{h\ell} + (1-\xi)(n_\ell/x_{h\ell})\, dx_{h\ell}$. Housing-firm FOCs imply $\xi p_\ell n_\ell/k_{h\ell} = r$ and $(1-\xi)p_\ell n_\ell/x_{h\ell} = q_\ell$, equivalently $\xi n_\ell/k_{h\ell} = r/p_\ell$ and $(1-\xi)n_\ell/x_{h\ell} = q_\ell/p_\ell$. Substituting and multiplying by p_ℓ,

$$p_\ell\, dn_\ell = r\, dk_{h\ell} + q_\ell\, dx_{h\ell}. \tag{25}$$

Substituting (23)–(24) into (22) and using (25) to collapse the housing-factor terms,

$$\begin{aligned}
\frac{1}{1+\lambda}\frac{dY}{dx_j} &= r\frac{dK}{dx_j} - r\sum_\ell \frac{dk_{h\ell}}{dx_j} + \sum_\ell w_\ell \frac{dn_\ell}{dx_j} + q_j - \sum_\ell q_\ell \frac{dx_{h\ell}}{dx_j} \\
&= q_j + r\frac{dK}{dx_j} - \sum_\ell p_\ell \frac{dn_\ell}{dx_j} + \sum_\ell w_\ell \frac{dn_\ell}{dx_j} \\
&= q_j + r\frac{dK}{dx_j} - \sum_\ell (p_\ell - w_\ell)\frac{dn_\ell}{dx_j},
\end{aligned} \tag{26}$$

which yields

$$\frac{dY}{dx_j} = (1+\lambda)\left[q_j + r\frac{dK}{dx_j} - \sum_\ell (p_\ell - w_\ell)\frac{dn_\ell}{dx_j}\right]. \tag{27}$$

We now combine (27) with the resource constraint. Substituting $dc/dx_j = dY/dx_j - \delta\, dK/dx_j$ into (20),

$$\frac{dW}{dx_j} = \frac{u_c}{1-\beta}\left[\frac{dY}{dx_j} - \delta\frac{dK}{dx_j} + \sum_\ell (p_\ell - w_\ell)\frac{dn_\ell}{dx_j}\right]. \tag{28}$$

Substituting (27) for dY/dx_j and collecting dK/dx_j and $\sum_\ell (p_\ell - w_\ell)\, dn_\ell/dx_j$ terms,

$$\frac{dW}{dx_j} = \frac{u_c}{1-\beta}\left[(1+\lambda)q_j + \big((1+\lambda)r - \delta\big)\frac{dK}{dx_j} + \big(1-(1+\lambda)\big)\sum_\ell (p_\ell - w_\ell)\frac{dn_\ell}{dx_j}\right], \tag{29}$$

which simplifies, using $1-(1+\lambda)=-\lambda$, to

$$\frac{dW}{dx_j}=\frac{u_c}{1-\beta}\left[(1+\lambda)q_j+((1+\lambda)r-\delta)\frac{dK}{dx_j}-\lambda\sum_\ell(p_\ell-w_\ell)\frac{dn_\ell}{dx_j}\right]. \tag{30}$$

The household Euler equation at the steady state gives $r+1-\delta=1/\beta$, hence $(1+\lambda)r-\delta=(1-\beta)/\beta+\lambda r$. Substituting and collecting terms,

$$\frac{dW}{dx_j}=\frac{u_c}{1-\beta}(1+\lambda)q_j+\frac{u_c}{\beta}\frac{dK}{dx_j}+\frac{\lambda\, u_c}{1-\beta}\left[r\frac{dK}{dx_j}-\sum_\ell(p_\ell-w_\ell)\frac{dn_\ell}{dx_j}\right]. \tag{31}$$

Finally, using (27) to substitute $r\,dK/dx_j-\sum_\ell(p_\ell-w_\ell)\,dn_\ell/dx_j=(dY/dx_j)/(1+\lambda)-q_j$ into the bracketed expression,

$$\frac{dW}{dx_j}=\frac{u_c}{1-\beta}[(1+\lambda)-\lambda]q_j+\frac{u_c}{\beta}\frac{dK}{dx_j}+\frac{\lambda}{1+\lambda}\cdot\frac{u_c}{1-\beta}\cdot\frac{dY}{dx_j}, \tag{32}$$

so that, in general,

$$\frac{dW}{dx_j}=\frac{u_c}{1-\beta}q_j+\frac{u_c}{\beta}\frac{dK}{dx_j}+\frac{\lambda}{1+\lambda}\cdot\frac{u_c}{1-\beta}\cdot\frac{dY}{dx_j}. \tag{33}$$

Evaluating at the pre-policy steady state ($\Delta x=0$, so $q_j=q_{j,0}$ and $c=c_0$), using $dW/d\Delta x_j=dW/dx_j$, and reinstating the "ss" subscript yields (11). □

B.2 Proof of Theorem 2

Proof of Theorem 2. Since $x_{jt}=x_{j,0}+\Delta x_j$, $dK_{ss}/d\Delta x_j=dK_{ss}/dx_j$; we derive the identity with respect to x_j. The firm's FOC for capital, $k_{y\ell}=\theta y_\ell/r$, combined with the production technology $y_\ell=A_\ell\bar{y}_\ell^{1+\lambda}$ where $\bar{y}_\ell=k_{y\ell}^\theta n_\ell^\chi(\alpha_\ell x_{y\ell})^\varphi$, and substituting y_ℓ out, yields

$$k_{y\ell}^{1-\theta(1+\lambda)}=(\theta A_\ell/r)\,n_\ell^{\chi(1+\lambda)}(\alpha_\ell x_{y\ell})^{\varphi(1+\lambda)}. \tag{34}$$

Under the condition $\theta(1+\lambda)<1$, this gives $k_{y\ell}\propto n_\ell^{\chi(1+\lambda)/(1-\theta(1+\lambda))}\,(\alpha_\ell x_{y\ell})^{\varphi(1+\lambda)/(1-\theta(1+\lambda))}$; direct differentiation yields

$$dk_{y\ell}=\frac{k_{y\ell}}{n_\ell}\cdot\frac{\chi(1+\lambda)}{1-\theta(1+\lambda)}dn_\ell+\frac{k_{y\ell}}{x_{y\ell}}\cdot\frac{\varphi(1+\lambda)}{1-\theta(1+\lambda)}dx_{y\ell}. \tag{35}$$

For housing, the agglomeration externality does not enter the production function, so inverting the binding feasibility $n_\ell=k_{h\ell}^\xi(\alpha_\ell x_{h\ell})^{1-\xi}$ gives $k_{h\ell}=n_\ell^{1/\xi}(\alpha_\ell x_{h\ell})^{-(1-\xi)/\xi}$, so

$$dk_{h\ell} = \frac{k_{h\ell}}{n_\ell} \cdot \frac{1}{\xi} \, dn_\ell - \frac{k_{h\ell}}{x_{h\ell}} \cdot \frac{1-\xi}{\xi} \, dx_{h\ell}. \tag{36}$$

Differentiating $K_{ss} = \sum_\ell (k_{y\ell} + k_{h\ell})$ with respect to x_j and grouping by driver using $\{dx_{y\ell}, dx_{h\ell}, dn_\ell\}$ yields (12). □

B.3 Proof of Theorem 3

Proof of Theorem 3. Problem (16) is a linear program in $\{\Delta x_j\}$, so KKT conditions are necessary and sufficient. The Lagrangian is

$$\mathcal{L} = \sum_j \frac{\partial W_{ss}}{\partial \Delta x_j}\bigg|_0 \Delta x_j - \mu\Big(\sum_j \frac{\kappa \, q_{j,0}}{1-\beta} \Delta x_j - B\Big) + \sum_j \eta_j \, \Delta x_j,$$

with raw multipliers $\mu \geq 0$ on the budget constraint and $\eta_j \geq 0$ on the non-negativity constraints. Stationarity in Δx_j gives

$$\frac{\partial W_{ss}}{\partial \Delta x_j}\bigg|_0 - \mu \, \frac{\kappa \, q_{j,0}}{1-\beta} + \eta_j = 0.$$

Dividing by $u_c \, \kappa \, q_{j,0}/(1-\beta) > 0$ and using the definition $\Pi_j^W \equiv (1-\beta)/(u_c \, \kappa \, q_{j,0}) \cdot (\partial W_{ss}/\partial \Delta x_j)|_0$ yields

$$\Pi_j^W - \frac{\mu}{u_c} + \frac{(1-\beta)\, \eta_j}{u_c \, \kappa \, q_{j,0}} = 0.$$

Defining the rescaled multipliers $\tilde{\mu} \equiv \mu / u_c$ and $\tilde{\eta}_j \equiv (1-\beta)\, \eta_j / (u_c \, \kappa \, q_{j,0})$ delivers the stated stationarity condition $\Pi_j^W - \tilde{\mu} + \tilde{\eta}_j = 0$. Complementary slackness $\eta_j \, \Delta x_j^* = 0$ and budget feasibility $\mu(\sum_j \kappa \, q_{j,0} \, \Delta x_j^* / (1-\beta) - B) = 0$ rescale to the stated forms $\tilde{\eta}_j \, \Delta x_j^* = 0$ and $\tilde{\mu}(B - \sum_j \kappa \, q_{j,0} \, \Delta x_j^* / (1-\beta)) = 0$.

For any two active regions j, ℓ with $\Delta x_j^*, \Delta x_\ell^* > 0$, complementary slackness gives $\tilde{\eta}_j = \tilde{\eta}_\ell = 0$, so the stationarity condition collapses to $\Pi_j^W = \Pi_\ell^W = \tilde{\mu}$. If a single region j^* strictly dominates all others in Π^W (i.e., $\Pi_{j^*}^W > \Pi_\ell^W$ for all $\ell \neq j^*$), the budget multiplier must satisfy $\tilde{\mu} = \Pi_{j^*}^W$, which forces $\tilde{\eta}_\ell = \tilde{\mu} - \Pi_\ell^W > 0$ for every $\ell \neq j^*$ and hence $\Delta x_\ell^* = 0$; the entire budget is allocated to j^*. □

C Analytical Supplements

This appendix collects auxiliary analytical results referenced in Section 3. Appendix C.1 records an equivalent organization of the welfare derivative decomposition (Theorem 1) that exposes the implicit dependence of the externality channel on the capital response. Appendix C.2 records the equivalent elasticity form of the capital response decomposition (Theorem 2), which is convenient for dimensionless comparative statics and for discrete-counterfactual accounting.

C.1 Alternative Welfare Derivative Decomposition

The welfare derivative decomposition of Theorem 1 is not unique. Since dY_{ss}, dK_{ss}, and $\{dn_\ell\}$ are all endogenous objects tied by the equilibrium conditions, any one of them can be substituted out in favor of the others, producing an equivalent but differently-organized expression. The following corollary records a useful alternative that exposes the implicit dependence of the externality channel on the capital response.

Corollary 1 (Alternative decomposition). *Using the identity $dY_{ss}/d\Delta x_j = (1+\lambda)\left[q_{j,0} + r\, dK_{ss}/d\Delta x_j - \sum_\ell (p_\ell - w_\ell)\, dn_\ell/d\Delta x_j\right]$ established in the proof of Theorem 1, the welfare derivative (11) can equivalently be written as*

$$\frac{dW_{ss}}{d\Delta x_j} = \underbrace{\frac{(1+\lambda)\, u_c\, q_{j,0}}{1-\beta}}_{\substack{\text{social direct} \\ \text{land channel}}} + \underbrace{u_c \left[\frac{1}{\beta} + \frac{\lambda\, r}{1-\beta}\right] \frac{dK_{ss}}{d\Delta x_j}}_{\substack{\text{capital accumulation channel} \\ \text{(externality-adjusted)}}} - \underbrace{\frac{\lambda\, u_c}{1-\beta} \sum_\ell (p_\ell - w_\ell)\, \frac{dn_\ell}{d\Delta x_j}}_{\substack{\text{labor reallocation channel} \\ \text{(externality-adjusted)}}}. \tag{37}$$

At $\lambda = 0$ the last term vanishes and the second reduces to $(u_c/\beta)\, dK_{ss}/d\Delta x_j$, recovering (11).

Proof. Substitute the $dY_{ss}/d\Delta x_j$ identity into the externality component of (11) and collect terms. □

Equations (11) and (37) reflect two equivalent ways of organizing the same welfare derivative. The form in Theorem 1, which we use throughout the body of the paper, isolates the *private* MP of land $q_{j,0}$ as the direct term and exposes the externality wedge as a single component proportional to $\lambda/(1+\lambda) \cdot dY_{ss}/d\Delta x_j$. The alternative form in (37) treats the *social* MP of land $(1+\lambda)q_{j,0}$ as the direct term and folds the externality wedge into the capital-accumulation and labor-reallocation channels, making explicit that the agglomeration force ultimately operates through the same general-equilibrium adjustment as the capital channel.

C.2 Elasticity Form of the Capital Response Decomposition

The level decomposition in Theorem 2 can be written in elasticity form. This representation is useful because it separates each channel into a capital-share weight and a dimensionless equilibrium elasticity.

Proposition 1 (Elasticity form of the capital response decomposition). *At the pre-policy steady state with $\theta(1+\lambda) < 1$, define*

$$\varepsilon_j^K \equiv \frac{x_j}{K_{ss}} \frac{dK_{ss}}{dx_j} = \frac{\partial \ln K_{ss}}{\partial \ln x_j}, \quad \varepsilon_{x_j}^{z_\ell} \equiv \frac{x_j}{z_\ell} \frac{dz_\ell}{dx_j} = \frac{\partial \ln z_\ell}{\partial \ln x_j},$$

for $z_\ell \in \{x_{y\ell}, x_{h\ell}, n_\ell\}$. Then the capital response decomposition can be written as

$$\varepsilon_j^K = \underbrace{\sum_\ell \omega_\ell^{xy} \varepsilon_{x_j}^{x_{y\ell}}}_{\substack{\text{land–capital complementarity} \\ \text{(production)}}} - \underbrace{\sum_\ell \omega_\ell^{xh} \varepsilon_{x_j}^{x_{h\ell}}}_{\substack{\text{land–capital substitution} \\ \text{(housing)}}} + \underbrace{\sum_\ell \omega_\ell^{\text{lab}} \varepsilon_{x_j}^{n_\ell}}_{\text{labor adjustment}}, \tag{38}$$

with weights

$$\omega_\ell^{xy} \equiv \frac{k_{y\ell}}{K_{ss}} \cdot \frac{\varphi(1+\lambda)}{1-\theta(1+\lambda)}, \quad \omega_\ell^{xh} \equiv \frac{k_{h\ell}}{K_{ss}} \cdot \frac{1-\xi}{\xi}, \quad \omega_\ell^{\text{lab}} \equiv \frac{k_{y\ell}}{K_{ss}} \cdot \frac{\chi(1+\lambda)}{1-\theta(1+\lambda)} + \frac{k_{h\ell}}{K_{ss}} \cdot \frac{1}{\xi}.$$

All objects are evaluated at the pre-policy steady state.

Proof. Multiplying (12) by x_j/K_{ss} gives

$$\varepsilon_j^K \equiv \frac{x_j}{K_{ss}} \frac{dK_{ss}}{dx_j} = \sum_\ell \frac{x_j}{K_{ss}} \Gamma_\ell^{xy} \frac{dx_{y\ell}}{dx_j} - \sum_\ell \frac{x_j}{K_{ss}} \Gamma_\ell^{xh} \frac{dx_{h\ell}}{dx_j} + \sum_\ell \frac{x_j}{K_{ss}} \Gamma_\ell^{\text{lab}} \frac{dn_\ell}{dx_j}.$$

The production-land term satisfies

$$\frac{x_j}{K_{ss}} \Gamma_\ell^{xy} \frac{dx_{y\ell}}{dx_j} = \frac{x_{y\ell}\Gamma_\ell^{xy}}{K_{ss}} \left(\frac{x_j}{x_{y\ell}} \frac{dx_{y\ell}}{dx_j} \right) = \frac{k_{y\ell}}{K_{ss}} \cdot \frac{\varphi(1+\lambda)}{1-\theta(1+\lambda)} \varepsilon_{x_j}^{x_{y\ell}}.$$

The housing-land term satisfies

$$\frac{x_j}{K_{ss}} \Gamma_\ell^{xh} \frac{dx_{h\ell}}{dx_j} = \frac{x_{h\ell}\Gamma_\ell^{xh}}{K_{ss}} \left(\frac{x_j}{x_{h\ell}} \frac{dx_{h\ell}}{dx_j} \right) = \frac{k_{h\ell}}{K_{ss}} \cdot \frac{1-\xi}{\xi} \varepsilon_{x_j}^{x_{h\ell}}.$$

The labor term satisfies

$$\frac{x_j}{K_{ss}} \Gamma_\ell^{\text{lab}} \frac{dn_\ell}{dx_j} = \frac{n_\ell \Gamma_\ell^{\text{lab}}}{K_{ss}} \left(\frac{x_j}{n_\ell} \frac{dn_\ell}{dx_j} \right) = \left[\frac{k_{y\ell}}{K_{ss}} \cdot \frac{\chi(1+\lambda)}{1-\theta(1+\lambda)} + \frac{k_{h\ell}}{K_{ss}} \cdot \frac{1}{\xi} \right] \varepsilon_{x_j}^{n_\ell}.$$

Substituting these three expressions into the first display yields (38), with the weights defined above. □

C.3 Ex Ante Decomposition and Finite Counterfactual Evaluation

The main counterfactual section uses two related but distinct objects. The allocation rule is based on the 2004 ex ante regional derivative (14). The aggregate outcomes in Table 5 evaluate finite allocations in the full nonlinear model. This appendix records how the policy-level decomposition tables aggregate the ex ante regional objects, and how the finite diagnostic accounting table should be interpreted.

Let p index a policy allocation, and let ω_j^p denote the model procurement-budget share assigned to region j under policy p:

$$\omega_j^p \equiv \frac{\kappa q_{j,2004} \Delta x_j^p / (1-\beta)}{\sum_\ell \kappa q_{\ell,2004} \Delta x_\ell^p / (1-\beta)}, \qquad \sum_j \omega_j^p = 1. \tag{39}$$

For the actual NTP row in Table 7, ω_j^p is the model procurement-budget share implied by the actual regional land increments, not the raw observed spending share. The policy-level ex ante return is the procurement-budget-weighted average of regional returns:

$$\Pi^W(p) = \sum_j \omega_j^p \Pi_j^W. \tag{40}$$

Writing the regional decomposition as

$$\Pi_j^W = \Pi_j^{W,\text{direct}} + \Pi_j^{W,\text{move}} + \Pi_j^{W,\text{ext}}, \tag{41}$$

the policy-level components are

$$\Pi^{W,m}(p) = \sum_j \omega_j^p \Pi_j^{W,m}, \qquad m \in \{\text{direct}, \text{move}, \text{ext}\}. \tag{42}$$

This is the aggregation used in Table 7.

The capital accumulation subdecomposition in Table 8 applies the same procurement-budget weighting after splitting each region's capital response. Let $D_j^K = dK_{\text{ss}}/d\Delta x_j$ and let $D_j^{K,r}$ be the contribution of subcomponent r in (12) and (13). Then

$$\Pi^{W,\text{move},r}(p) = \sum_j \omega_j^p \, \Pi_j^{W,\text{move}} \frac{D_j^{K,r}}{D_j^K}, \qquad r \in \{xy, xh, \tilde{n}, N\}. \tag{43}$$

The four subcomponents in (43) sum to $\Pi^{W,\text{move}}(p)$ up to rounding.

The equality of the direct term in Table 7 is a statement about normalized ex ante returns. From (14), $\Pi_j^{W,\text{direct}} = 1/\kappa$ for every region because the raw direct derivative is divided by a procurement cost proportional to the same pre-policy rent. The raw regional direct derivative itself is not equal across regions:

$$\frac{dW_{\text{ss}}^{\text{direct}}}{d\Delta x_j} = \frac{u_c}{1-\beta} q_{j,0}. \tag{44}$$

Thus high-rent regions such as Seoul can have much larger raw direct marginal returns, even though their normalized direct returns are the same. In the diagnostic finite accounting table below, the equal direct entries for the two same-budget policies are a separate aggregate accounting fact: when the same model procurement budget is spent, $\sum_j q_{j,0}\Delta x_j^p$ is fixed by construction. It does not imply that raw regional direct derivatives are equal.

For completeness, the diagnostic finite accounting formula used for the appendix table is

$$\Delta W_p^{\text{diag}} = \underbrace{\frac{u_c^b}{1-\beta}\sum_j q_j^b \Delta x_j^p}_{\text{direct}} + \underbrace{\frac{u_c^b}{\beta}\left(K_p - K_b\right)}_{\text{capital accumulation}} + \underbrace{\frac{\lambda}{1+\lambda}\frac{u_c^b}{1-\beta}\left(Y_p - Y_b\right)}_{\text{externality}} + \varepsilon_p, \tag{45}$$

where b denotes the no-NTP benchmark and ε_p collects the finite-change residual. This diagnostic accounting is not the allocation rule; the allocation rule uses the 2004 ex ante derivative aggregation above.

Table C.1: Diagnostic Accounting for Finite Counterfactuals

Comparison	Total welfare change	Direct MP land	Capital accumulation	Externality	Residual
Actual NTP	0.070	0.056	0.013	0.002	-0.001
(ratio to total)	(1.00)	(0.79)	(0.19)	(0.03)	(-0.01)
Ex ante optimal NTP	0.088	0.056	0.027	0.003	0.002
(ratio to total)	(1.00)	(0.63)	(0.31)	(0.04)	(0.03)

Notes: The table applies (45) to finite steady-state differences relative to No NTP. Entries are model accounting units, not consumption-equivalent percentages in Table 5. This table is included as a diagnostic and is not used to rank the 2004-informed allocation.

D The 2nd New Town Project and Counterfactuals

This appendix describes how we map the 2nd NTP into the land-supply counterfactuals. The goal is to measure the land supplied by the actual project and its development cost, then use those objects to construct the no-NTP and same-budget ex ante optimal NTP experiments.

The 2nd NTP consisted of eleven subprojects. Nine were located in the SMA and account for nearly 90% of the supplied land. Two non-SMA subprojects were located in Daejeon and Chungcheongnam-do (D&C). Our analysis use the SMA subprojects to measure the actual land increment removed in the no-NTP counterfactual. The D&C subprojects provide useful cost information because they were part of the same policy wave but were developed outside the high-cost SMA land market.

Our measure of urban land is a sum of residential, commercial, and industrial land, matching the land input used in the quantitative model. The project data report land supplied by use and total expenditure, but not use-specific expenditure. We therefore allocate total expenditure across land-use categories using each category's land-supplied share.

In the no-NTP counterfactual, we remove only the actual SMA land supplied by the 2nd NTP. The D&C projects are treated as part of the existing 2019 land stock rather than being removed. This convention keeps the no-NTP experiment focused on the main object of the paper: the large-scale supply of new urban land to the SMA and its implications for regional concentration. The same convention is used when comparing the actual NTP to the optimal same-budget allocation.