Making room on an empty planet: The spatial consequences of depopulation*

Tim de Silva
Stanford GSB

James D. Paron
Stanford GSB

June 17, 2026

Latest Version

Abstract

We study the effects of population decline on the spatial distribution of economic activity. Empirically, we document that population growth is negatively associated with changes in spatial concentration, both across and within countries. We then show theoretically that an aggregate population decline raises spatial concentration when housing supply elasticities decrease or expenditure shares increase with population density across regions, both of which are constant in benchmark spatial models. To quantify this mechanism, we estimate a dynamic spatial model using Japanese data and regional variation in housing supply elasticities and expenditure shares. In response to the projected population decline, the model predicts that Tokyo's population share rises from 10% to 40%, while low-density regions empty out. A social planner would like to increase concentration further by making transfers from low- to high-density regions.

* First draft: June 2026. We are grateful to Erica Bucchieri for outstanding research assistance. We also thank Adrien Auclert, Adrien Bilal, Jesús Fernández-Villaverde, Chad Jones, Benny Kleinman, Arvind Krishnamurthy, Hanno Lustig, Monika Piazzesi, Steve Redding, Martin Schneider, Chris Tonetti, and Stijn Van Nieuwerburgh for helpful comments and suggestions. Emails: tdesilva@stanford.edu and jparon@stanford.edu.

Fertility is below replacement in many countries, meaning populations are expected to decline rapidly (Geruso and Spears 2026). In some countries, this decline has already begun. For example, Japan's population peaked in 2010 and is projected to fall by 50% by 2100 (Gietel-Basten 2016). A declining population will likely have significant aggregate consequences, affecting economic growth (Jones 2022; Peters and Walsh 2026), fiscal sustainability (Samuelson 1958; Aaron 1966), capital accumulation (Baker et al. 2005; Weil 2026), and the environment (Ehrlich and Holdren 1971). But given that economic activity is unevenly distributed across space (Redding and Rossi-Hansberg 2017), the consequences of depopulation may not be the same everywhere. Indeed, concerns about the uneven spatial consequences of depopulation have led several governments to introduce policies that are targeted at rural areas.[1] In this paper, we ask: how does depopulation affect the spatial allocation of economic activity, and what are the implications for optimal spatial policy?

The main takeaway of this paper is that depopulation increases spatial concentration. The mechanism we propose operates through housing. Because land is in fixed supply, housing is a scarce resource that forces people to spread out. When the total population of a country declines, housing becomes less scarce everywhere, but especially so in ex-ante dense areas. In equilibrium, this drives migration into dense areas and increases spatial concentration.

We establish this finding in four steps. First, we present empirical evidence of *scale dependence*: within and across countries, changes in national population are negatively correlated with changes in spatial concentration. Second, we show that in a wide class of spatial equilibrium models, scale dependence arises if ex-ante denser regions have lower housing supply elasticities or higher housing expenditure shares, both of which are consistent with empirical evidence. Third, we quantify the importance of this mechanism by estimating a dynamic spatial model using Japanese data, along with the existing regional estimates of housing supply elasticities and expenditure shares. Our main quantitative result is that Japan's projected population decline leads to a large rise in spatial concentration, with Tokyo's population share quadrupling in the long run. Finally, we show that a social planner would like to increase this concentration further by making transfers from low- to high-density regions.

In the first part of the paper, we present empirical evidence that population growth is negatively associated with changes in spatial concentration. We begin with a case study of Japan, where the population level has been declining since 2010. During the decades in which total population was increasing, there was no discernible increase in spatial concentration; in fact, the population share of Tokyo (the densest region) fell. In contrast, once Japan's population began to decline, spatial

[1] Examples include local *akiya* banks and renovation grants for abandoned homes in Japan; "1 Euro home" programs tied to houses in remote Italian and Spanish regions; and the Addressing Depopulation Action Plan in Scotland, which provides various sources of funding to rural areas.

concentration began to rise, as people in sparsely populated regions (e.g., Hokkaido) migrated to densely populated regions (e.g., Tokyo).

The scale dependence in Japan is present in a large panel of countries. Using population estimates from the Global Human Settlement Layer, we test whether national population changes are associated with changes in spatial concentration, defined as the share of the population living in high-density areas. Across countries, we find a strong negative relationship: countries that experienced larger population growth saw greater declines in spatial concentration. This same pattern also holds within countries over time, as in Japan: spatial concentration rises (falls) when population falls (rises). These relationships are quantitatively unchanged when controlling for other drivers of urbanization, such as structural transformation (Michaels et al. 2012; Coeurdacier et al. 2025), population aging (Giannone et al. 2026; Badilla Maroto et al. 2026), and GDP growth (Eckert et al. 2025).[2]

Motivated by this empirical evidence, we next characterize the conditions under which scale dependence arises in spatial equilibrium. We start with a standard static equilibrium model of location choice with housing. Households consume both goods and housing services, where total housing supply is a composite of land (a fixed factor) and structures (a variable factor, scaling with local labor). As usual, productivity differences and agglomeration benefits concentrate households, while scarce housing supply and taste shocks spread them out. Our first result is that equilibrium population shares are invariant to the population level (scale *independence*) if and only if both housing supply and housing demand are Cobb-Douglas.[3] This is the prevailing assumption in spatial models (e.g., Helpman 1998; Allen and Arkolakis 2014; Ahlfeldt et al. 2015; Caliendo et al. 2019; Kleinman et al. 2023), which in turn predict that depopulation has no spatial consequences.

Our second result is that, for general specifications of housing supply and demand, population shares depend on the population level when housing supply elasticities, demand elasticities, or expenditure shares vary across regions. This clarifies why Cobb-Douglas generates scale independence: with Cobb-Douglas, these objects are constant and equal everywhere. To generate a negative relationship between population and spatial concentration—as in the data—we need that regional housing supply elasticities decrease with density or that expenditure shares increase with density.[4] Both of these conditions are supported by evidence from existing literature. In particular, using estimates of region-level housing supply elasticities from Baum-Snow and Han (2024), we show that

[2] In Section 1.2, we discuss how the negative relationship between population growth and spatial concentration is consistent with the longer history of urbanization. This history reflects two offsetting forces: population growth, which reduces concentration; and structural transformation out of agriculture, which increases concentration.

[3] We focus on housing because of the empirical evidence on housing supply elasticities and expenditure shares. However, in several extensions, we show our theory applies in complementary ways to goods and amenities.

[4] We do not consider the case in which price elasticities of housing demand vary across regions because, to our knowledge, there are no regional estimates of these elasticities. Additionally, the commonly used models of housing demand imply that this elasticity is constant (Cobb-Douglas and PIGL) or approximately constant (CES).

supply elasticities decline steeply with population density. On the demand side, we summarize the evidence from Combes et al. (2019) and Finlay and Williams (2025), which suggests that housing expenditure shares increase with density (Robert-Nicoud et al. 2026).

This regional variation in housing supply and demand that drives scale dependence arises naturally from *complementarity* (i.e., elasticities of substitution less than one). Consider first the supply side. If structures and land are complements in production, then supply elasticities will be lower in high-density places like Tokyo. The intuition is that land is a "weak link" in Tokyo: structures are abundant, but per-capita housing supply is constrained by the scarce supply of land. In contrast, land is abundant in sparse places like Hokkaido, so that adding more people does almost nothing to per-capita housing supply. Consequently, when the population falls, housing supply rises disproportionately in Tokyo as land becomes freed up, incentivizing migration.

A similar intuition applies on the demand side when goods and housing are complements in consumption. In high-density places like Tokyo, goods are abundant, but that abundance is of limited value given the relative scarcity of housing. The marginal utility of housing is thus high in Tokyo, so residents spend a higher share of their budgets on housing. Therefore, when population falls, the resulting increase in housing supply is valued more in Tokyo. This incentivizes people to move to ex-ante denser places, increasing spatial concentration.

Next, we quantify this mechanism by studying its implications for the projected depopulation in Japan. To do so, we develop a dynamic spatial equilibrium model in which housing supply elasticities and expenditure shares endogenously vary across regions. Households decide where to live and work, subject to moving frictions (Caliendo et al. 2019; Kleinman et al. 2023). A competitive construction sector owns the land and maintains a durable capital stock of structures (Glaeser and Gyourko 2005). We deviate from Cobb-Douglas on the supply side by modeling housing as a CES aggregate of land and structures; on the demand side, we use PIGL preferences (Muellbauer 1975, 1976). We estimate the model and invert its fundamentals from Japanese data in 2000. Importantly, our estimation includes the CES and PIGL parameters that govern complementarity in housing supply and demand, which we identify using the regional variation in supply elasticities and expenditure shares described above. We then analyze the transition path for Japan's projected depopulation from around 125M people in 2000 to 10M people in 2195.

Our main quantitative result—from which all others follow—is that spatial concentration dramatically increases. In the long run, the population share of Tokyo quadruples from around 10% to 40%. The population shares of the next two densest regions also increase. In contrast, the remaining, less dense regions experience significant population declines. An extreme case is Hokkaido, whose share falls from over 4% to 0.1%. Around half of these changes are due to complementarity in housing on

the supply side (CES production) and half from complementarity in demand (PIGL demand). In contrast, with Cobb-Douglas supply and demand, population shares would not change, and Tokyo's population would fall from 12M to 1M (versus 4M in the estimated model).

Due to agglomeration benefits, changes in per-capita incomes (i.e., wages) mirror changes in population levels and fall everywhere in the long run. However, because of the housing-driven migration response, wages fall much less in Tokyo than in Hokkaido (5% versus 25%). Despite declining local wages, aggregate GDP per capita rises 27% in the long run because of the migration to higher-productivity places like Tokyo (versus a 12% fall with scale independence).

Importantly, spatial concentration begins to rise well in anticipation of the population decline. Although we do not start the aggregate population decline until 2020, Tokyo's population level rises by 3% between 2000 and 2020, while Hokkaido's falls by 3%. This anticipation reflects the fact that households are subject to moving frictions, so they choose to move to Tokyo in anticipation of future housing supply growth. Quantitatively, these changes in population shares account for around one-quarter of observed changes without any concurrent changes in model primitives or differential birth rates. This anticipatory migration also drives house prices up in Tokyo and down in Hokkaido. However, once aggregate population declines, the rise in housing supply is so large that house prices eventually fall everywhere.

In the final part of the paper, we study how optimal spatial policy responds to depopulation. We have two main results. First, at the initial population level, a utilitarian social planner wants to transfer resources away from sparse regions (Hokkaido) and toward dense regions (Tokyo). Second, when the aggregate population falls, the planner wants to increase transfers towards Tokyo even more. As in Fajgelbaum and Gaubert (2020), optimal transfers balance an agglomeration externality, which calls for taxing Hokkaido, with redistribution, which calls for taxing Tokyo. However, in our dynamic model, moving costs create an additional force: because households cannot fully arbitrage utility differences across regions, the planner can raise welfare by encouraging migration toward regions with higher average utility. Scale dependence amplifies this force because housing becomes disproportionately less scarce in dense regions as population falls, making Tokyo relatively more attractive and leading the planner to increase transfers toward Tokyo and away from Hokkaido. Importantly, these are statements about steady state transfers; along the transition, the planner may want to move more slowly to protect households who are temporarily locked into declining regions.

Related literature. This paper builds on the large literature on economic geography, which extends the spatial equilibrium models of Rosen (1979) and Roback (1982) to quantify how various forces shape allocations and optimal policies across and within space (see Redding and Rossi-Hansberg 2017; Allen and Arkolakis 2025; Redding 2025; Fajgelbaum and Gaubert 2025; Desmet and Parro

2025; Diamond and Suárez Serrato 2025 for reviews). We make three contributions to this literature. First, we establish general conditions for scale dependence within this class of models, highlighting the role of complementarity in housing supply and demand. Second, we integrate a rich dynamic model of housing into a quantitative model and show how housing generates substantial changes in the spatial distribution of economic activity in response to a population decline. Third, we show how optimal spatial policy responds to depopulation, which has important implications for the growing number of spatially targeted policies that have been designed in response to depopulation.

To generate scale dependence, we deviate from Cobb-Douglas housing supply and demand, which are the prevailing assumptions in existing spatial models. Importantly, we discipline these deviations using evidence from several studies in urban economics.[5] On the supply side, estimates of housing supply elasticities are negatively related to regional density (Saiz 2010; Baum-Snow and Han 2024) and estimates of the land share of housing value are increasing in density (Davis and Heathcote 2007; Davis et al. 2021). We explain these facts by modeling housing as a CES aggregate of land and structures with substitution elasticity less than one.[6] On the demand side, several papers provide evidence that regional housing expenditure shares are increasing in population density (Combes et al. 2019; Finlay and Williams 2025). Following existing literature that attempts to explain these facts (Gaubert and Robert-Nicoud 2025; Robert-Nicoud et al. 2026), we use price-independent generalized linear (PIGL) preferences (Muellbauer 1975, 1976), which nest Cobb-Douglas and have a constant price elasticity of housing demand.[7] We use PIGL for convenience, but show that CES demand with complements would achieve the same result.

We also contribute to the literature studying the economic consequences of population decline (Samuelson 1958; Aaron 1966; Ehrlich and Holdren 1971; Baker et al. 2005; Jones 2022; Peters and Walsh 2026). This literature has mainly focused on aggregate consequences, whereas our contribution is to highlight spatial consequences. Importantly, the focus of this paper is on the consequences of population level independently of the effects of demographic changes and aging. The aggregate effects of demographic change have also been studied extensively, with less focus on spatial consequences (e.g., Auerbach et al. 1989; Mankiw and Weil 1989; Cutler et al. 1990; Poterba 2001; Bloom et al. 2010; Acemoglu and Restrepo 2017; Auclert et al. 2021). While demographics may be important, we think it is fruitful to isolate the spatial consequences of population level for

[5] An elasticity of substitution less than one for both housing supply and demand is consistent with Rognlie (2015), who argues that both are needed to explain the large increase in the housing sector's capital income share.

[6] An alternative model of housing supply that could also fit these facts is Cobb-Douglas production and heterogeneous fixed costs of land development (Combes et al. 2021). While this model would also generate scale dependence, in Section 3 we discuss why we prefer CES production.

[7] In general, PIGL allows for a price elasticity and income elasticity that both differ from one. Because we are interested in price effects, we set the income elasticity to one (i.e., homothetic utility) in our baseline model, but explore the consequences of non-homotheticities in extensions.

two reasons. First, empirically, we find that the relationship between spatial concentration and population levels is similar even if we condition on age distributions. Second, in the long run, the effects of aging are likely limited because, assuming fertility and mortality rates stabilize, the age distribution converges to a stable (i.e., constant) distribution even as the population level continues to decline (Weil 2026). Nevertheless, population aging can have important spatial consequences as fertility continues to decline. For example, Giannone et al. (2026) show that the aging population in Japan has led to a decline in the amenities of sparse areas due to the out-migration of (more mobile) young workers. We view this mechanism as complementary and important, especially for helping to explain what has happened in Japan from 2000 to present, and show in an extension that endogenous amenities can amplify our main housing mechanism.[8]

Finally, a growing literature investigates the underlying causes of the fertility decline (see Doepke et al. (2023) and Gobbi et al. (2026) for reviews). One strand of this literature uses spatial models to study the drivers of falling fertility, finding that rising costs of education, childcare, and unpriced urban congestion in large cities are contributing factors (Arkolakis et al. 2026; Borck et al. 2026). Another strand has argued that fertility responds to housing supply and, therefore, that rising housing costs can explain part of the fertility decline (Lovenheim and Mumford 2013; Dettling and Kearney 2014; Fazio et al. 2025; Couillard 2026). While our main results take the population decline as given, we consider an extension in which fertility responds to housing. We show that the predictions for long-run fertility can be dramatically different depending on whether housing supply is Cobb-Douglas. With Cobb-Douglas, fertility always bounces back to replacement because per-capita housing increases without bound. In contrast, with complements, fertility may remain permanently below replacement because the increase in supply is bounded.

1 Empirical evidence of scale dependence

In this section, we present empirical evidence of *scale dependence*: population growth is negatively associated with spatial concentration. We first present evidence using administrative data from Japan, where the aggregate population has been declining and will likely continue given a total fertility rate that is well below replacement. We then show that these patterns within Japan generalize to a large panel of countries.

[8] Badilla Maroto et al. (2026) present complementary evidence that retirees are more able to arbitrage spatial differences and migrate to poor rural regions, given that they no longer participate in the labor market. In turn, this migration has positive effects on local economic activity, especially in the services sector.

Figure 1. Aggregate population and total fertility rate in Japan

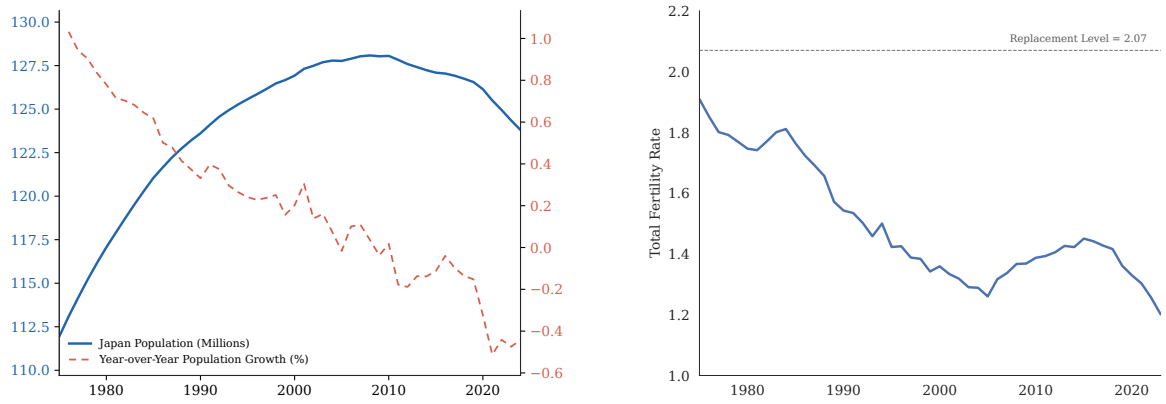

Notes: The left panel shows the total population from the Statistics Bureau of Japan; the dashed series is the annual percent change in that total. The right panel plots Japan's official annual total fertility rate from the Ministry of Health, Labour and Welfare Vital Statistics final-report workbook; the dashed line is the replacement level of 2.07 reported by the National Institute of Population and Social Security Research.

1.1 A case study within Japan

Japan is the canonical example of a country with a declining population and is therefore a natural place to begin looking for evidence of scale dependence. The left panel of Figure 1 shows that Japan's total population grew from 112 million in 1975 to a peak of 128 million in 2010, then began declining such that in 2020 the population was around the same level as it was in 2000. The right panel shows the main cause: Japan's total fertility rate (TFR) has been below replacement since the early 1970s, averaging 1.35 from 2000 to 2023. Absent a significant increase in TFR, Japan's population will continue to decline rapidly in the coming years (Giannone et al. 2026).

Data. We use annual data on population growth and land area by region from the Statistics Bureau of Japan between 1975–2020. We divide Japan into its 47 prefectures, which are geographic regions that represent the highest level of administrative division below the national government.[9] For each prefecture j and decade $t \in \{1980, 1990, 2000, 2010\}$, we compute (i) the change in the log of j's population share between t and $t + 10$ and (ii) the log of j's population share at the start of the decade over its fixed 2010 inhabitable land share, a normalized measure of initial density. We use the log rather than level of each variable because the model in subsequent sections predicts this relationship is approximately linear in logs, though results are similar in levels.

Results. Figure 2 shows a scatterplot of changes in prefecture population shares against each prefecture's initial population density for each decade, where the size of each dot is proportional

[9] See Figure A.1 for a map of Japan's prefectures.

to the initial population of the prefecture. Tokyo and Hokkaido—Japan's densest and sparsest prefectures, respectively—are highlighted in red. Prior to Japan's population peak (1980–1990 and 1990–2000), the relationship between these variables is relatively weak. In contrast, after the population approaches its peak and begins declining (2000–2010 and 2010–2020), the initial density and subsequent population growth become more positively correlated. Quantitatively, Tokyo's population share rose by 1.6 percentage points between 2000 and 2019, while Hokkaido's fell by 0.3 percentage points—magnitudes that we will attempt to explain in our dynamic model.[10] The fact that spatial concentration increases around the time that the population begins to decline provides suggestive evidence of scale dependence.[11]

Before turning to our cross-country evidence, we discuss three additional results that are helpful for interpreting the increase in spatial concentration from Figure 2. First, Figure A.2 shows that a large fraction of the increase in Tokyo's population share reflects internal migration rather than differences in birth or death rates, consistent with Giannone et al. (2026). Second, Figure A.3 documents that this increase in migration is concentrated among young people in their twenties, who account for the bulk of in-migrants in every year, a fact emphasized by Giannone et al. (2026) and Mori and Murakami (2025). Finally, Figure A.4 shows that the migration into Tokyo is relatively evenly distributed across origin prefectures. The post-2010 increase in concentration is therefore not a feature of one or two specific origin regions but a broad-based reallocation of population toward denser locations.

1.2 Systematic evidence across and within countries

Our case study in Japan provides suggestive evidence that a decline in the aggregate population leads to increased spatial concentration. We now ask whether the same pattern appears systematically in a large set of countries.

Data. To measure spatial concentration systematically across countries, we use data from the Global Human Settlement Layer (GHSL) population grid, which provides population counts on roughly $1\,\text{km}^2$ grids over the entire world. The GHSL is produced by the European Commission's Joint Research Centre and combines census data, satellite imagery, and UN population data to construct consistent, high-resolution population counts (Schiavina et al. 2023).[12] The GHSL data are available

[10] Giannone et al. (2026) use similar data to document a related fact in Japan: municipalities, which are subdivisions of prefectures, with older (younger) populations have seen population decreases (increases) since 1980.

[11] Interestingly, Glaeser et al. (2025) document a positive relationship between initial population and subsequent population growth across cities in Ukraine (Figure 2, Panel B), whose population has been declining since the collapse of the Soviet Union in 1991. They also report that Kyiv's population grew while most large Ukrainian cities shrank.

[12] For details on the construction of GHSL, see Pesaresi et al. (2024) and the GHSL Data Package.

Figure 2. Initial population density and subsequent population growth: Japan

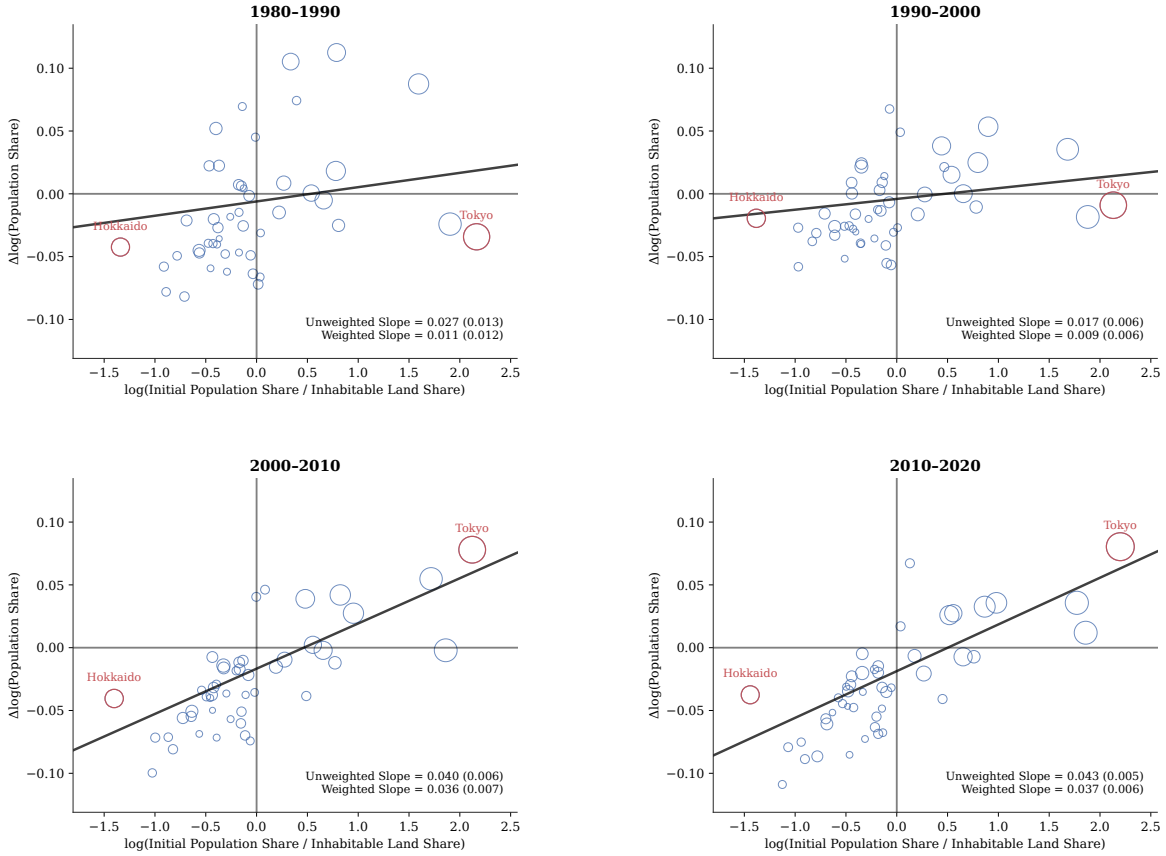

Notes: The figure plots the relationship between initial population density and subsequent population-share growth across Japanese prefectures. We use annual prefecture population and inhabitable land shares from the Statistics Bureau of Japan. Each marker represents one prefecture, with marker areas proportional to initial population. The horizontal axis is the log of initial population share divided by inhabitable land share, and the vertical axis is the change in log population share over the decade. The black fitted line is the initial-population-weighted regression line; the legend reports unweighted and initial-population-weighted slopes with heteroskedasticity-robust standard errors in parentheses.

in five-year intervals, and we start our sample in 2000.[13] We aggregate cells to country boundaries from the Global Administrative Unit Layers and restrict the sample to countries that can be matched to ISO identifiers, excluding standard overseas territories and a small set of small-country exclusions.

For each country, we summarize concentration with a single statistic: the share of the country's population that lives in cells classified as "high-density" as of 2000. To define this set, we (i) sort cells by density within each country, (ii) find the cell-population cutoff in 2000 such that approximately 20% of the country's population lives in cells above it, (iii) fix this set of cells, and (iv) compute the share of the country's population living in this fixed set in 2020.[14] The two key variables for each

[13]We choose to start our sample in 2000 because this is when the GHSL begins using GPWv4 to compute its population estimates. Before 2000, we could not find clear documentation on how GHSL's estimates are constructed.

[14]Fixing the high-density set in 2000 ensures that changes in the share over time reflect changes in where people live, not changes in which cells are classified as high-density.

Figure 3. Population growth and spatial concentration across countries

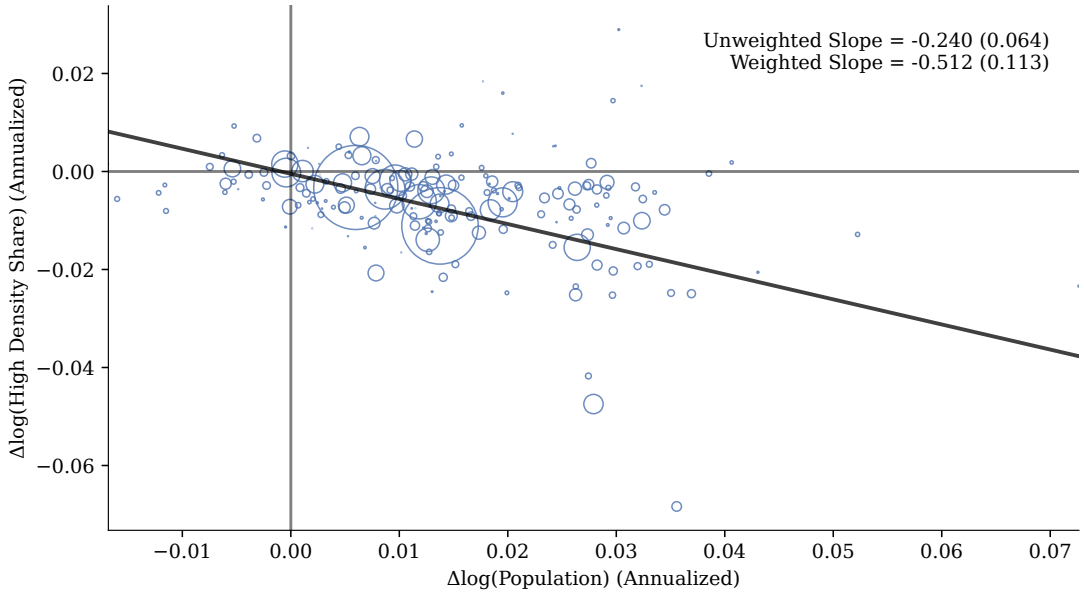

Notes: This figure shows the relationship between annualized population growth and annualized changes in spatial concentration across countries between 2000 and 2020. Each marker is a country, where the size of each circle is proportional to the total population in that country in 2000. The fitted line is a 2000-population-weighted regression line; the legend reports unweighted and 2000-population-weighted slopes with heteroskedasticity-robust standard errors in parentheses.

country are the annualized log change in this high-density share between 2000 and 2020 and the annualized log change in the country's total population over the same window.[15] We also collect each country's agricultural employment share, GDP, and old-age dependency ratio (population 65 and older relative to population 15–64) from the World Bank World Development Indicators. We construct GDP per capita by dividing GDP by the GHSL total population count.

Results. Figure 3 shows a scatterplot of the relationship between population growth and the change in the high-density share across countries, where the size of each circle is proportional to the total population in that country in 2000. The slope is negative, indicating that countries with slower population growth saw larger increases in the share of their population concentrated in dense areas. Quantitatively, a 1pp reduction in population growth is associated with an increase in the annualized growth of the high-density share of approximately 0.2–0.5pp.

Table 1 reports the corresponding regression results to Figure 3 with several different specifications. We perform both unweighted and population-weighted regressions, in each case adding two controls. The first control is the change in a country's agricultural employment share, which addresses the possibility that the increase in concentration and urban sprawl is driven by structural transformation

[15] In Figure A.5, we verify the quality of the GHSL data by showing that the population counts for Japanese prefectures are nearly perfectly correlated with the administrative data from the previous section.

Table 1. Cross-country evidence of scale dependence

	$\Delta \log$(High Density Share)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta \log$(Population)	-0.240	-0.197	-0.215	-0.178	-0.512	-0.394	-0.466	-0.384
	(0.064)	(0.057)	(0.074)	(0.072)	(0.113)	(0.112)	(0.125)	(0.120)
$\Delta \log$(Agriculture Share)		-0.120		-0.117		-0.133		-0.130
		(0.070)		(0.071)		(0.065)		(0.070)
$\Delta \log$(Old-Age Dependency Ratio)			0.047	0.037			0.051	0.013
			(0.064)	(0.072)			(0.057)	(0.071)
Constant	-0.003	-0.006	-0.004	-0.006	-0.000	-0.005	-0.002	-0.005
	(0.001)	(0.002)	(0.001)	(0.002)	(0.001)	(0.002)	(0.002)	(0.003)
2000 Population Weights	No	No	No	No	Yes	Yes	Yes	Yes
Observations	190	178	190	178	190	178	190	178
R^2	0.093	0.124	0.096	0.126	0.301	0.330	0.303	0.330

Notes: The table reports cross-country regressions that characterize the relationship between population growth and changes in spatial concentration between 2000 and 2020. All variables measured as changes are annualized. Columns (1)–(4) are unweighted; columns (5)–(8) are weighted by 2000 population. Standard errors are shown in parentheses and are heteroskedasticity-robust.

out of agriculture into services and manufacturing (Michaels et al. 2012; Coeurdacier et al. 2025). The second control is the old-age dependency ratio. We include this control because our interest is in the effects of the population *level* on spatial concentration rather than changes in the age distribution of the population. Aging could affect mobility if younger individuals are more mobile (Greenwood 1997; Lucas 1997; Giannone et al. 2026), retirees are more able to arbitrage spatial differences (Badilla Maroto et al. 2026), or because aging drives structural transformation (Cravino et al. 2022). Across all eight specifications in Table 1, the slope on population growth is negative and similar in magnitude to the unconditional estimates in Figure 3.

Robustness. We conduct several robustness checks on our findings in Figure 3 and Table 1. First, the negative relationship holds both within and across continents (Figure A.6, Table A.1). Second, the relationship between population growth and spatial concentration is quantitatively similar when we control for the growth in GDP per capita in each country, which may drive urbanization to the extent that growth is urban-biased (Eckert et al. 2025) (Table A.2). Third, the pattern is not an artifact of the log specification: measuring concentration using raw rather than log changes delivers similar patterns (Table A.3). Finally, the pattern is similar when we construct the analogue of Figure 3 by alternatively defining high-density regions as the top 0.1% densest cells in 2000 within each country (Figure A.7).

Panel evidence. While the previous evidence exploits variation across countries, we now exploit variation within countries over time, similar to the evidence from Japan in Section 1.1. For each country i and five-year window starting in $t \in \{2000, 2005, 2010, 2015\}$—the GHSL periods we use—we compute the annualized log change in the high-density share between t and $t+5$, keeping

Table 2. Within-country evidence of scale dependence

	$\Delta \log$(High Density Share)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\Delta \log$(Population)	-0.244	-0.241	-0.255	-0.244	-0.229	-0.248	-0.216
	(0.045)	(0.045)	(0.129)	(0.130)	(0.133)	(0.126)	(0.131)
$\Delta \log$(Agriculture Share)					0.005		0.004
					(0.018)		(0.017)
$\Delta \log$(Old-Age Dependency Ratio)						0.011	0.020
						(0.027)	(0.028)
Constant	-0.003	-0.002	0.001	0.001	0.000	0.001	-0.000
	(0.001)	(0.001)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Country Fixed Effects	No	No	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	No	Yes	No	Yes	No	No	No
Observations	760	760	760	760	712	760	712
R^2	0.080	0.084	0.643	0.647	0.643	0.643	0.643

Notes: The table reports panel regressions that characterize the relationship between population growth and changes in spatial concentration over five-year intervals between 2000 and 2020. All variables measured as changes are annualized. Columns include different combinations of fixed effects. Standard errors are heteroskedasticity-robust.

fixed the cells classified as high-density in 2000 using the procedure described above, as well as the corresponding annualized log changes in the total population, agricultural employment share, and old-age dependency ratio. Table 2 shows results from the analogous regressions to Table 1 in these panel data.

The first two columns of Table 2 show that the relationship between concentration and population growth is similar across all the five-year intervals to the estimate in Table 1. The third and fourth columns introduce country fixed effects, exploiting within-country variation in population growth. While the standard error naturally increases given the short panel, the coefficients are quantitatively very similar to the first two columns and to the estimates in Table 1. These results suggest that the connection between population growth and concentration that we highlighted within Japan in Section 1.1 generalizes to a larger set of countries. The final three columns add changes in agricultural employment shares and old-age dependency ratios as controls. While these controls reduce the coefficient on population growth slightly, their estimated coefficients are imprecise and differ more substantially in magnitude from Table 1.

Longer historical context. The negative relationship between population growth and spatial concentration that we document is not inconsistent with the longer history in which population growth and urbanization occurred together. We interpret this history as reflecting two offsetting forces: population growth, which we find reduces concentration; and structural transformation out of agriculture, which existing literature argues increases concentration (Michaels et al. 2012; Coeurdacier et al. 2025). As in the broader structural transformation literature (Herrendorf et al. 2014), the latter force requires one of two historical trends: growth in the relative productivity of

agricultural goods with inelastic demand across sectors or a decline in the preference for agriculture due to non-homothetic preferences. Either of these trends will cause the conversion of agricultural land into non-agricultural land, which increases spatial concentration in non-agricultural areas because agricultural production is more land-intensive.[16] Historically, it is plausible that structural transformation was the dominant force, given that agricultural shares fell substantially.[17] In contrast, structural transformation will likely be a weaker force moving forward, both because many countries already have very low and stable agricultural shares and because a declining population would, if anything, slow down technological innovation (Jones 2022; Peters and Walsh 2026).

2 A theory of depopulation and concentration

The empirical evidence points to *scale dependence* in spatial concentration: concentration falls (rises) when the total level of population rises (falls). We now lay out a framework for thinking about the conditions needed to explain scale dependence in spatial equilibrium.

2.1 Definition of scale dependence

There is a mass N of households, each of which chooses from a set of locations $i \in \{1, \ldots, J\}$ to maximize utility. This gives rise to equilibrium population shares N_i/N. Assume that these shares are positive and unique.[18] We define scale (in)dependence as follows.

Definition (Scale (in)dependence)**.** *An equilibrium exhibits* ***scale independence*** *if population shares N_i/N are invariant to total population N. That is, it is scale-independent if*

$$\frac{d(N_i/N)}{dN} = 0, \quad \forall i \in \{1, \ldots, J\}.$$

Otherwise, we say that the equilibrium exhibits ***scale dependence****.*

Notice that scale independence is equivalent to having $d \log N_i / d \log N = d \log N_j / d \log N$ for all locations i and j. This will prove to be a more convenient condition to work with.

2.2 Intuition in a simple static model

We begin with the simplest possible setting to capture the essential intuition. We will progressively generalize this setting in each subsequent section.

[16] In Appendix C.3 (discussed in Section 2.5), we formalize this result in an extension of our model with multiple goods.

[17] For example, in the U.S., the population grew from 76 million in 1900 to 281 million in 2000, while the agricultural share fell from 40% to 2%. Given these changes, the estimates in Table 1 suggest concentration would decrease.

[18] In all of the cases that follow, we will assume that the parameters give rise to a unique, interior equilibrium. Allen et al. (2024) give a thorough analysis of these conditions. The assumption of positive population shares is not strictly necessary, but allows us to ignore uninteresting corner solutions.

Environment. There is a population N of households, indexed $k \in [0, N]$. Each household chooses a location $i_k \in \{0, \ldots, J\}$, goods consumption c_{ik}, and housing services h_{ik} to solve

$$V_k = \max_{i, c_{ik}, h_{ik}} c_{ik}^{1-\alpha} h_{ik}^{\alpha} \epsilon_{ik}, \quad \alpha \in (0, 1), \tag{1}$$

$$w_i \geq c_{ik} + p_i h_{ik}, \tag{2}$$

where w_i is the wage and p_i is the rental price of housing.[19] The taste shocks ϵ_{ik} are i.i.d. across individuals and locations and are drawn from a Fréchet distribution,

$$\mathbb{P}(\epsilon_{ik} \leq \epsilon) = \exp\left\{-\epsilon^{-\frac{1}{\theta}}\right\}, \quad \epsilon > 0. \tag{3}$$

The parameter $\theta \geq 0$ governs the dispersion of the taste shocks: as $\theta \downarrow 0$, there is no dispersion; as $\theta \to \infty$, there is infinite dispersion. Throughout this subsection, assume $\theta > (1 - \alpha)\sigma$, so that equilibrium population shares are non-degenerate.

Goods are produced locally by firms, with total output given by the production function

$$Y_i = (A_i N_i^{\sigma}) N_i, \quad \sigma > 0. \tag{4}$$

The parameter σ captures external increasing returns to scale from agglomeration (Krugman 1991). The constants A_i are primitive, location-specific productivities capturing innate geographic advantages of some locations over others. We assume all firms earn zero profits.

Finally, the total supply of housing is a composite $H_i^{\mathrm{S}}(X_i, L_i)$ of structures X_i and land L_i.[20] The quantity of structures scales with the quantity of local labor, N_i, as $X_i = A_X N_i$. The quantity of land L_i is fixed. This implies per-capita local housing supply

$$h_i^{\mathrm{S}}(N_i) \equiv \frac{H_i^{\mathrm{S}}(X_i, L_i)}{N_i}, \quad X_i = A_X N_i.$$

While it may at first seem strange to assume that housing supply is exogenous, not responding to prices p_i or wages w_i, this simplistic setting will prove to be quite close to the result that arises from profit-maximizing firms, which we will consider shortly.

[19] Here and elsewhere, we assume goods are undifferentiated and trade is costless across regions, so we can normalize the price to one everywhere. The free-trade assumption is inconsequential, because households end up consuming what they produce; but it will matter when we consider spatial transfers in Section 6.

[20] We will assume $H_i^{\mathrm{S}}(X_i, L_i)$ is twice-differentiable in X_i for mathematical convenience.

Equilibrium. Conditional on choosing location i, (1)–(2) imply constant expenditure shares

$$c_i = (1-\alpha)w_i \quad \text{and} \quad p_i h_i = \alpha w_i.$$

Notice that choices are symmetric within i (we have dropped k subscripts). The optimal location choices are then

$$i_k \in \arg\max_j \left\{c_j^{1-\alpha} h_j^{\alpha} \epsilon_{jk}\right\} = \arg\max_j \left\{w_j^{1-\alpha} h_j^{\mathsf{S}}(N_j)^{\alpha} \epsilon_{jk}\right\},$$

where the second equality imposes housing-market clearing—demand $h_i^{\mathsf{D}}(w_i, p_i)$ equals supply $h_i^{\mathsf{S}}(N_i)$. By the properties of the Fréchet distribution (3), this implies that the share of the total population N living in location i solves

$$\frac{N_i}{N} = \frac{(w_i^{1-\alpha} h_i^{\mathsf{S}}(N_i)^{\alpha})^{\frac{1}{\theta}}}{\sum_j (w_j^{1-\alpha} h_j^{\mathsf{S}}(N_j)^{\alpha})^{\frac{1}{\theta}}}. \tag{5}$$

We can simplify this further by using (4) to substitute in the goods-sector wage $w_i = A_i N_i^{\sigma}$:

$$\frac{N_i}{N} = \frac{(A_i^{1-\alpha} h_i^{\mathsf{S}}(N_i)^{\alpha})^{\frac{1}{\theta-(1-\alpha)\sigma}}}{\sum_j (A_j^{1-\alpha} h_j^{\mathsf{S}}(N_j)^{\alpha})^{\frac{1}{\theta-(1-\alpha)\sigma}}}. \tag{6}$$

Notice that, even though w_i depends on scale N_i through the agglomeration force, it does not create scale dependence because the elasticities σ are homogeneous across locations. It does, however, increase the level of spatial concentration through the exponent $1/(\theta-(1-\alpha)\sigma)$.[21]

Conditions for scale dependence. It is immediately apparent from looking at the population shares (6) that scale dependence hinges on the functional form of housing supply $h_i^{\mathsf{S}}(N_i)$. To see how, define the supply elasticities (with respect to population decline),

$$\varepsilon_i^h(N_i) \equiv -\frac{\partial \log h_i^{\mathsf{S}}(N_i)}{\partial \log N_i}, \tag{7}$$

and differentiate (6) with respect to N to get the population responses

$$\frac{d\log N_i}{d\log N} = \frac{\alpha}{\theta-(1-\alpha)\sigma} \frac{\partial \log h_i^{\mathsf{S}}(N_i)}{\partial \log N_i} \frac{d\log N_i}{d\log N} + \tilde{\Gamma}$$

[21] We choose not to focus on heterogeneous agglomeration elasticities as a source of scale dependence for two reasons. First, such scale dependence would likely be in the wrong direction because, if anything, agglomeration elasticities are likely higher in cities (see, e.g., Davis and Dingel 2019). However, (5) shows that a positive correlation between σ_i and N_i would predict that spatial concentration would decrease as the population falls. Second, (8) suggests that agglomeration elasticities are likely a smaller source of scale dependence than housing: Ahlfeldt and Pietrostefani (2019) survey existing estimates of σ and report a mean of 0.05, which is an order of magnitude smaller than typical estimates of housing supply elasticities (Saiz 2010; Baum-Snow and Han 2024).

$$= \frac{1}{\theta - (1-\alpha)\sigma + \alpha\varepsilon_i^h(N_i)}\Gamma, \tag{8}$$

where $\tilde{\Gamma}$ and Γ are normalizing constants, identical for all i. Hence, there is scale dependence if and only if there are heterogeneous supply elasticities. We formalize this as follows.

Proposition 1 (Scale independence in the simple model)**.** *The simple static model exhibits scale independence if and only if housing supply elasticities are identical: $\varepsilon_i^h(N_i) = \varepsilon_j^h(N_j)$ for all i and j. These elasticities are identical for all primitives if and only if housing supply takes the form*[22]

$$H_i^S(X_i, L_i) = G_i(L_i)X_i^{1-\nu} \tag{9}$$

for some functions $G_i(L_i)$ and some constant ν.

The first statement of the proposition follows directly from (8). The second statement—that scale independence is equivalent to (9)—means that $\varepsilon_i^h(N_i)$ is equal to a constant ν everywhere and at all population levels.[23] The typical example of (9) is the Cobb-Douglas aggregator $H_i^S(X_i, L_i) = \bar{H}_i X_i^{1-\nu} L_i^{\nu}$, $\nu \in [0, 1]$. Cobb-Douglas housing is thus scale-independent.

This statement may seem knife-edge. However, many—if not most—existing regional and urban spatial models satisfy (9), and therefore exhibit scale independence (e.g., Helpman 1998; Redding 2016; Allen and Arkolakis 2014; Ahlfeldt et al. 2015; Caliendo et al. 2019; Kleinman et al. 2023).[24] Thus, a large subset of existing spatial frameworks predict that population decline will have no spatial consequences.

Finally, consider the direction of scale dependence: for which locations does a decline in N cause an increase in population share? Take two locations i and j. From (8), we can see that a decline in N causes N_i/N_j to rise if and only if $\varepsilon_i^h(N_i) > \varepsilon_j^h(N_j)$. The location with the largest housing supply elasticity $\varepsilon_i^h(N_i)$ sees the largest increase in per-capita housing supply in response to a population decline, and thus sees an increase in its population share N_i/N. The reverse is true for the smallest-elasticity location. We next impose some more structure on housing to make clearer what determines differences in these elasticities.

[22] "All primitives" means all combinations of parameters admitting a unique interior equilibrium.

[23] Together with $X_i = A_X N_i$, (9) implies $h_i^S(N_i) = \bar{h}_i N_i^{-\nu}$ for constants $\bar{h}_i$ independent of N_i.

[24] In our model, housing is a non-tradable good, but the implications for scale dependence are identical for any (priced or non-priced) amenity, like congestion or crime (Diamond 2016). What is relevant is not that housing is sold and consumed, but that it enters utility in proportion to $N_i^{-\alpha\nu}$.

A tractable special case: CES housing supply. An intuitive and tractable functional form for housing supply, which we will assume throughout the paper, is the CES aggregator

$$H_i^S(X_i, L_i) = \left((1-\nu) X_i^{\frac{\eta-1}{\eta}} + \nu L_i^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}}, \quad \nu \in (0,1). \tag{10}$$

If $\eta = 1$, we return to the Cobb-Douglas (hence scale-independent) case above. If instead $\eta < 1$, then structures and land are complements: building many structures has diminishing value if land supply remains low. In Section 3, we discuss in detail why we choose this model of housing production and provide evidence that $\eta < 1$ is the empirically plausible case.

This aggregator implies that per-capita housing supply—the amount each resident gets to consume—is

$$h_i^S(N_i) = \left((1-\nu) A_X^{\frac{\eta-1}{\eta}} + \nu \left(\frac{L_i}{N_i} \right)^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}},$$

and hence the supply elasticities (7) are

$$\varepsilon_i^h(N_i) = \frac{\nu (A_X N_i / L_i)^{\frac{1-\eta}{\eta}}}{(1-\nu) + \nu (A_X N_i / L_i)^{\frac{1-\eta}{\eta}}} \in [0, 1).$$

For $\eta < 1$, the elasticity is increasing with local population density N_i/L_i. That is, lowering the population disproportionately increases per-capita housing supply in high-density areas when land and structures are complements. From (8), that means that lowering the aggregate population N causes the population share to rise in the highest-density region (e.g., Tokyo) and fall in the lowest-density region (e.g., Hokkaido).

As we did in the data in Section 1, define *spatial concentration* as the share of population in the (ex-ante) densest $\bar{J} < J$ places. Assume no "ties," so the high-density set is unique.

Proposition 2 (Scale independence with CES housing supply)**.** *Assume housing supply takes the CES form (10). The equilibrium is scale-independent for all primitives if and only if $\eta = 1$ (Cobb-Douglas). Moreover, a decline in population N causes an increase in spatial concentration if $\eta < 1$ (complements), and a decrease in concentration if $\eta > 1$ (substitutes).*

In the complements case, land becomes a "weak link" in high-density regions: at some point, building additional structures no longer increases total housing supply, so the only effect of adding more people is to further divvy up the scarce supply of land. In contrast, when density is low, there is a dearth of structures and an abundance of land, so adding more people (and structures) has little effect on per-capita housing supply. Now flip this logic in the other direction: a population *decline* disproportionately raises housing supply in Tokyo by freeing up land, but has no effect in Hokkaido,

so people are incentivized to migrate to Tokyo.

2.3 General conditions for scale dependence

The simple static model delivers most of the intuition. Still, it is unrealistic to assume that labor and housing supply are exogenous to market prices. Additionally, we have focused on scale effects in housing supply, but there may just as well be demand-side effects. Here, we generalize the model along these (and other) dimensions to derive some more general results.

Environment. Each household continues to solve (1)–(3), but we now generalize the utility aggregator from Cobb-Douglas ($c_{ik}^{1-\alpha}h_{ik}^{\alpha}$) to a general form $U(c_{ik}, h_{ik})$. We assume that U is homogeneous of degree one, but this is relaxed in the appendix.

Local labor supply (N_i) is now allocated between goods production (N_{iY}) and construction of structures (N_{iX}). Labor is perfectly substitutable between these tasks, so the wage w_i is common to both sectors.[25] Goods are produced according to[26]

$$Y_i = (A_i N_i^{\sigma}) N_{iY}, \quad \sigma > 0. \tag{11}$$

Housing $H_i(X_i, L_i)$ is built and rented out by absentee construction firms who own the underlying land.[27] Structures are built with labor according to $X_i = A_X N_{iX}$ and fully depreciate between periods. Thus, the construction sector maximizes profits

$$\Pi_i = \max_{N_{iX}} \{p_i H_i(X_i, L_i) - w_i N_{iX}\}, \quad X_i = A_X N_{iX}. \tag{12}$$

We will abuse notation and write $H_i(X_i)$, except when land L_i is important. To ensure a well-behaved solution to (12), assume $H_i'(X_i) > 0$, $H_i''(X_i) < 0$, $\lim_{X_i \downarrow 0} H_i'(X_i) = \infty$, and $\lim_{X_i \to \infty} H_i'(X_i) = 0$.

Equilibrium. The maximization of direct utility $U(c_i, h_i)$ implies an indirect utility function $U^*(w_i, p_i)$. Since U is homogeneous, $U^*(w_i, p_i) = w_i v(p_i)$ for some decreasing function $v(p_i)$. By Roy's identity, the Marshallian demand for housing is then given by

$$h_i^{\mathrm{D}}(w_i, p_i) = -\frac{w_i}{p_i} \frac{\partial \log v(p_i)}{\partial \log p_i}, \tag{13}$$

[25] We will relax this in the dynamic model, which allows for occupational choice and hence wage differences.

[26] Results are similar with sector-level agglomeration, i.e., N_{iY}^{σ} instead of N_i^{σ}. We assume this in Section 4.

[27] The fact that landlords are absentee means that profits leave the economy in the form of lost goods ($Y_i = C_i + \Pi_i$). This only becomes relevant for optimal policy (Section 6), so we address it in that section.

which implies $c_i^{\mathrm{D}}(w_i, p_i) = w_i - p_i h_i^{\mathrm{D}}(w_i, p_i)$. With Fréchet shocks, population shares solve

$$\frac{N_i}{N} = \frac{U(c_i, h_i)^{\frac{1}{\theta}}}{\sum_j U(c_j, h_j)^{\frac{1}{\theta}}} = \frac{U^*(w_i, p_i)^{\frac{1}{\theta}}}{\sum_j U^*(w_j, p_j)^{\frac{1}{\theta}}} = \frac{(A_i v(p_i))^{\frac{1}{\theta-\sigma}}}{\sum_j (A_j v(p_j))^{\frac{1}{\theta-\sigma}}}, \tag{14}$$

where the last equality substitutes in the goods sector's first-order condition $w_i = A_i N_i^{\sigma}$. Households prefer to live in locations with high productivity A_i and low house prices p_i.

Given the wage $w_i = A_i N_i^{\sigma}$, the construction sector's problem (12) determines the employment allocation $N_{iX} = N_i - N_{iY}$ via the first-order condition

$$w_i = p_i H_i'(X_i) A_X. \tag{15}$$

This condition gives rise to a housing supply curve $H_i^{\mathrm{S}}(p_i, w_i)$, implying per-capita supply

$$h_i^{\mathrm{S}}(w_i, p_i, N_i) \equiv \frac{H_i^{\mathrm{S}}(p_i, w_i)}{N_i}.$$

The market clears at $h_i^{\mathrm{S}}(w_i, p_i, N_i) = h_i^{\mathrm{D}}(w_i, p_i)$.

Conditions for scale dependence. As in the simple model, we can characterize scale dependence by comparing $d \log N_i / d \log N$ across locations. Differentiating (14) gives

$$\begin{aligned} \frac{d \log N_i}{d \log N} &= \frac{1}{\theta - \sigma} \frac{\partial \log v(p_i)}{\partial \log p_i} \frac{d \log p_i}{d \log N_i} \frac{d \log N_i}{d \log N} + \tilde{\Gamma} \\ &= \left(\theta - \sigma + \frac{p_i h_i}{w_i} \frac{d \log p_i}{d \log N_i} \right)^{-1} \Gamma, \end{aligned} \tag{16}$$

where $\tilde{\Gamma}$ and Γ are common normalizing constants and where the second equality substitutes in the housing demand curve (13). We now see two potential sources of scale dependence: heterogeneous housing expenditure shares $p_i h_i / w_i$ and heterogeneous local house-price responses $d \log p_i / d \log N_i$. In the simple model above, Cobb-Douglas utility meant $p_i h_i / w_i = \alpha$.

For a proportionate decline in house prices $d \log p_i / d \log N_i$, locations with higher housing expenditure shares will see a relative increase in population. This is because the expenditure share is a sufficient statistic for the marginal utility of housing (see (13)). That is, places with high expenditure shares have high marginal utility, so alleviating scarcity in high-expenditure regions raises utility disproportionately more, incentivizing people to move there. We will show below that this is naturally implied by housing and goods being complements in demand.

Of course, house prices are equilibrium objects, so their response to a change in N depends on simultaneous supply and demand responses. Totally differentiating $h_i^{\mathrm{S}}(w_i, p_i, N_i) = h_i^{\mathrm{D}}(w_i, p_i)$ and

the other equilibrium conditions with respect to N_i yields the following result.

Proposition 3 (Scale dependence in the general model)**.** *In the general static model, the response of local population to a change in total population is, up to a normalizing constant,*

$$\frac{d\log N_i}{d\log N} \propto \frac{1}{\theta - \sigma + \frac{p_i h_i}{w_i}\left(\frac{1}{\zeta_i^S + \zeta_i^D} + \frac{\zeta_i^S + 1}{\zeta_i^S + \zeta_i^D}\sigma\right)}, \tag{17}$$

where the housing supply and demand elasticities are given, respectively, by

$$\zeta_i^S \equiv \frac{\partial \log h_i^S}{\partial \log p_i} = -\frac{H_i'(X_i)}{H_i(X_i)}\frac{H_i'(X_i)}{H_i''(X_i)} > 0, \tag{18}$$

and

$$\zeta_i^D \equiv -\frac{\partial \log h_i^D}{\partial \log p_i} = -p_i\left(\frac{v''(p_i)}{v'(p_i)} - \frac{v'(p_i)}{v(p_i)}\right) > 0.$$

The equilibrium thus exhibits scale independence for all primitives if and only if

$$U(c_i, h_i) \propto c_i^{1-\alpha} h_i^{\alpha} \quad \textit{and} \quad H_i(X_i, L_i) = G_i(L_i) X_i^{1-\nu} \tag{19}$$

for some functions $G_i(L_i)$ and some constants $\alpha \in (0,1)$ and $\nu \in (0,1)$.

We learn two things from Proposition 3 relative to Proposition 1. First, even when housing supply responds to prices, the supply-side condition for scale independence in (19) comes down to technology $H_i(X_i, L_i)$. As before, a Cobb-Douglas housing technology implies scale independence. Second, we learn that Cobb-Douglas utility is not just sufficient, but necessary for scale independence. Hence, one can get scale dependence in two ways: either deviate from Cobb-Douglas housing demand or deviate from Cobb-Douglas housing supply.

As before, it will be most useful to think about the effect of population on spatial concentration after imposing more structure in the model. We do this next.

A tractable special case: CES housing supply and PIGL housing demand. We again assume CES housing supply (10). We now deviate from Cobb-Douglas on the demand side by assuming a particular form of non-unit-elastic housing demand.

First, consider the implications of CES housing supply. The supply elasticities (18) are

$$\zeta_i^S = \frac{1-\nu}{\nu}\eta\left(\frac{L_i}{X_i}\right)^{\frac{1-\eta}{\eta}}. \tag{20}$$

For $\eta = 1$, elasticities are constant; for $\eta < 1$, they are falling in the density of structures per unit of

land X_i/L_i. Because structures grow with the number of people, this will mean elasticities tend to fall with population densities. That is, housing supply is less responsive to prices in Tokyo (where supply is constrained by land) than in Hokkaido. Ultimately, the intuition for this is no different than it was in the simple static model. Land in Tokyo is a weak link.

To generalize housing demand, we assume the homothetic form of price-independent generalized linear (PIGL) preferences (Muellbauer 1975, 1976), for which indirect utility equals

$$U^*(w_i, p_i) = w_i \exp\left\{-\alpha \frac{p_i^{1-\zeta_p} - 1}{1-\zeta_p}\right\}. \tag{21}$$

In Section 3, we discuss in detail why we choose this model of housing demand and provide empirical evidence in support of it. Under PIGL, Marshallian demand (13) takes the simple form $h_i^{\mathrm{D}}(w_i, p_i) = \alpha w_i p_i^{-\zeta_p}$, implying constant demand elasticities and variable housing expenditure shares:

$$\zeta_i^{\mathrm{D}} = \zeta_p \quad \text{and} \quad \frac{p_i h_i}{w_i} = \alpha p_i^{1-\zeta_p}. \tag{22}$$

When $\zeta_p = 1$, PIGL collapses to Cobb-Douglas utility. When $\zeta_p < 1$, which Section 3 shows is the empirically relevant case, housing expenditure shares are higher in locations with high house prices. Intuitively, these will be the denser places, like Tokyo, in which limited housing supply and higher wages push up house prices.

We can again think of this in terms of complements and substitutes. The parameter ζ_p governs the (Hicksian) elasticity of substitution between housing and consumption:[28]

$$-\frac{\partial \log(h_i/c_i)}{\partial \log p_i} = \frac{\zeta_p - p_i h_i/w_i}{1 - p_i h_i/w_i}.$$

This elasticity is less than one (complements) if $\zeta_p < 1$ and greater than one (substitutes) if $\zeta_p > 1$. This is what determines whether expenditure shares rise or fall with prices.[29] Just like on the supply side, complementarity on the demand side means that housing is the weak link in high-priced regions. Goods are abundant in Tokyo, but the scarcity of housing makes it harder to enjoy them.

[28] We assume throughout that $\zeta_p > p_i h_i/w_i$, so this elasticity is positive.

[29] These inequalities therefore generate the same predictions about the direction of scale independence under both PIGL and CES, so the economic intuition is the same under either assumption. To be precise, assume the CES aggregator $U(c,h) = ((1-\alpha)c^{\frac{\mu-1}{\mu}} + \alpha h^{\frac{\mu-1}{\mu}})^{\frac{\mu}{\mu-1}}$. If the substitution elasticity $\mu < 1$ (complements), then both expenditure shares $p_i h_i/w_i = \alpha^\mu p_i^{1-\mu}/((1-\alpha)^\mu + \alpha^\mu p_i^{1-\mu})$ and demand elasticities $\zeta_i^{\mathrm{D}} = \mu + (1-\mu)(p_i h_i/w_i)$ increase in price p_i. It is then straightforward to show that PIGL with $\zeta_p < 1$ is qualitatively the same as having goods and housing be CES with $\mu < 1$, both of which correspond to the complements case.

Under CES supply and PIGL demand, the population changes (17) become

$$\frac{d\log N_i}{d\log N} \propto \left(\theta - \sigma + \alpha p_i^{1-\zeta_p}\left(\frac{1}{\zeta_i^S + \zeta_p} + \frac{\zeta_i^S + 1}{\zeta_i^S + \zeta_p}\sigma\right)\right)^{-1}. \tag{23}$$

Suppose that $\eta < 1$ and $\zeta_p < 1$, and that house prices p_i and structure density X_i/L_i are monotonically increasing in population density N_i/L_i.[30] Then (17) says that a decline in aggregate population will increase the population share in ex-ante high-density regions—people move from Hokkaido to Tokyo. On the supply side, $\eta < 1$ means that per-capita housing supply goes up disproportionately in dense areas; on the demand side, $\zeta_p < 1$ means that the marginal utility benefit of the higher housing supply (lower house prices) is larger in dense areas. Thus, spatial concentration rises when population falls.

Proposition 4 (Scale independence with CES supply and PIGL demand)**.** *Assume CES housing supply (10) and PIGL demand (21). The equilibrium is scale-independent for all primitives if and only if $\eta = \zeta_p = 1$ (Cobb-Douglas supply and demand). Moreover, if regions are the same size ($L_i = L$), then a decline in total population N causes an increase in spatial concentration if $\eta \leq 1$ (complements in supply) and $\zeta_p \leq 1$ (complements in demand), with at least one inequality strict.*

This is the key result for understanding this paper. While we will add numerous bells and whistles to the model, the economics comes down to this: with the conventional Cobb-Douglas supply and demand, population decline has no spatial consequences. The effects we do find will therefore all hinge on deviations from the Cobb-Douglas benchmark.

2.4 Wage and house price responses to population decline

What happens to local wages and house prices in response to depopulation? Wages are proportional to N_i^σ, and so

$$\frac{d\log w_i}{d\log N} = \sigma\frac{d\log N_i}{d\log N}. \tag{24}$$

Thus, the local wage response is proportional to the local population change.

The consequences for house prices are more nuanced. Applying the chain rule and (17) gives

$$\frac{d\log p_i}{d\log N} = \frac{d\log p_i}{d\log N_i}\cdot\frac{d\log N_i}{d\log N} \propto \frac{\left(\frac{1}{\zeta_i^S + \zeta_i^D} + \frac{\zeta_i^S + 1}{\zeta_i^S + \zeta_i^D}\sigma\right)}{\theta - \sigma + \frac{p_i h_i}{w_i}\left(\frac{1}{\zeta_i^S + \zeta_i^D} + \frac{\zeta_i^S + 1}{\zeta_i^S + \zeta_i^D}\sigma\right)}, \tag{25}$$

[30] A sufficient (but not necessary) condition for this to be true is to define locations with equal land sizes $L_i = L$, like we did with the GHSL grid data in Section 1.

In response to a population decline, house prices fall everywhere. Which places experience the largest declines, however, is ambiguous. On the one hand, if supply elasticities ζ_i^S fall in density (CES supply with $\eta < 1$), then prices will fall by more in denser areas.[31] This is because housing supply rises more in locations in which land is scarce. On the other hand, if expenditure shares $p_i h_i / w_i$ rise in density (PIGL demand with $\zeta_p < 1$), then prices fall by more in sparse areas. This is because of migration: population falls, increasing housing supply everywhere, so people move to dense areas, reducing relative supply in those areas and pushing up relative prices. We will quantify these two channels in Section 5.1.

2.5 Extensions

Before proceeding to the full dynamic model, we briefly discuss three extensions exploring the generality of these results. Details are relegated to Appendix C.

Trade, commuting, and amenities. All of the propositions in this section are unaffected by the following generalizations. First, one can allow for differentiated goods across locations (à la Armington) with iceberg trade costs. Second, one can allow for commuting across regions with iceberg commuting costs. And third, one can allow for amenities—including endogenous public goods or spillovers from congestion and crime—provided that they enter utility in the power form $a_i N_i^\gamma$ for some constants a_i and γ (the typical assumption in the literature).

Non-homothetic utility. We assumed that utility is homogeneous of degree one. Appendix C.2 re-derives all results with this assumption relaxed. With non-homothetic preferences, differences in income levels across regions can generate scale dependence. Under standard parameter values, these income effects would dampen the rise in spatial concentration in response to a population decline. However, as discussed in Appendix C.2, this dampening is extremely small quantitatively because its magnitude depends on σ, which is small. Therefore, we choose not to focus on it because it introduces additional complications, such as a non-trivial choice of the numeraire.

Scale effects in goods and amenities. We have focused on scale effects in housing (with some anticipation of the empirical evidence we will show later), but in principle there could also be scale effects in goods and amenities. Nothing in the theory above hinges on the label "housing." In Appendices C.3 and C.4, we consider two extensions, both of which produce qualitatively similar results to the housing model above:

1. **Multiple goods:** Following the structural transformation literature (Herrendorf et al. 2014), assume c_i is a composite of two goods s, agriculture ($s = a$) and non-agriculture ($s = n$).

[31] This comparative static requires $\theta > \sigma$, which is easily satisfied in our dynamic model.

Following Michaels et al. (2012), these goods are complements and are produced using labor and land as $Y_{is} = A_{is} N_{is}^{1-\mu_s} L_i^{\mu_s}$, with agriculture being more land-intensive ($\mu_a > \mu_n$). This implies two results. First, high-density regions specialize in non-agriculture. Second, a decline in aggregate population N causes people to move to high-density regions and work in non-agriculture, increasing spatial concentration.

2. **Multiple amenities:** Each local government optimally taxes its residents to provide two complementary types of amenities. Divisible amenities scale with per-capita spending. Shared-capacity amenities, such as hospitals, transit systems, libraries, or emergency-response networks, can serve many residents before congestion fully binds, so their cost per resident falls with local population. When N declines, low-density places have a harder time sustaining the shared-capacity part of the amenity bundle. Because that part complements divisible inputs, households migrate to high-density areas, increasing spatial concentration.

These examples illustrate a general theoretical takeaway. In order to get depopulation to cause a rise in spatial concentration, we need two things: complements and differential returns to scale in N_i. With CES housing, structures (which scale with N_i) were complementary to land (independent of N_i). With PIGL, goods (which scale with N_i^σ) were complementary to housing (which scales less than one-for-one). The multiple-goods and multiple-amenities examples we just summarized amplify spatial concentration in the same way.

3 Empirical evidence on housing supply and demand

The key result of Section 2 is equation (17) in Proposition 3, which states that population changes are

$$\frac{d \log N_i}{d \log N} \propto \frac{1}{\theta - \sigma + \dfrac{p_i h_i}{w_i} \left(\dfrac{1}{\zeta_i^{\mathrm{S}} + \zeta_i^{\mathrm{D}}} + \dfrac{\zeta_i^{\mathrm{S}} + 1}{\zeta_i^{\mathrm{S}} + \zeta_i^{\mathrm{D}}} \sigma \right)}.$$

This equation says that concentration increases when the population level falls—as documented in Section 1—if at least one of the following is true: (i) housing supply elasticities, ζ_i^{S}, are falling in population density N_i/L_i; (ii) housing expenditure shares, $p_i h_i / w_i$, are rising in N_i/L_i; or (iii) housing demand elasticities, ζ_i^{D}, are falling in N_i/L_i.

The purpose of this section is twofold. First, we provide and discuss empirical evidence that suggests conditions (i) and (ii) are likely to hold, while (iii) is likely to be quantitatively unimportant. Second, we explain why CES housing production and PIGL housing demand are sensible model ingredients for explaining this evidence.

3.1 Housing supply elasticities and CES production

We use estimates of housing unit supply elasticities from Baum-Snow and Han (2024) (BSH) between 2000 and 2010, which to our knowledge represent the state-of-the-art estimates of supply elasticities.[32] Given our baseline model does not have commuting, the ideal definition of a region is one in which there is a high degree of overlap between workplace and residential location. Therefore, we use BSH's region-level estimates for Census Transportation Planning Package (CTPP) regions, which are based on Census residence, workplace, and commuting-flow tabulations and roughly correspond to metropolitan labor-market areas. We then measure 2000 region population levels and density using the 2000 Census, assigning each Census tract to its corresponding CTPP regions.

The left panel of Figure 4 shows that housing supply elasticities are decreasing in the region's population density, providing direct evidence for condition (i) above that is needed to generate scale dependence. This finding is not completely surprising given the evidence in Saiz (2010) and BSH that big cities tend to have less elastic housing supply. In order to generate this fact, we use a CES production function for housing with $\eta < 1$. To understand why, notice that (20) implies[33]

$$\frac{d \log \zeta_i^S}{d \log(N_i/L_i)} \approx \frac{\eta - 1}{\eta}. \tag{26}$$

The relationship depicted in the left panel of Figure 4 is a semi-elasticity, but the corresponding elasticity is -0.29, implying $\eta \approx 0.77$. While the previous equation will not be exact in our dynamic model, the core idea is that $\eta < 1$ generates housing supply elasticities that decrease with density. This complementarity between structures and land is also consistent with Figure A.8, which shows that the land share of property values increases with density across regions (Davis and Heathcote 2007; Davis et al. 2021).[34]

An alternative model of housing supply that also generates housing supply elasticities that decrease with density is Cobb-Douglas production and heterogeneous fixed costs of land development (Combes et al. 2021; Baum-Snow and Han 2024). In this model, heterogeneity in housing supply elasticities comes entirely from differences in the distribution of fixed costs; conditional on land development, housing supply elasticities are constant. While this model is natural when housing production is driven by regulatory constraints, we prefer CES production in this context for two reasons. First, the empirical evidence in Saiz (2010) suggests that geography explains significantly more of the

[32] Ideally, we would use estimates from Japan, but we do not know of existing estimates of similar quality.

[33] This relationship is an approximation because X_i/N_i will vary across regions. In our dynamic model in Section 4, this variation is relatively small.

[34] Baum-Snow and Duranton (2025) survey the extensive literature that estimates the elasticity of substitution between land and capital in housing production. The range of estimates are extremely wide, but to the extent that estimates are different from one (i.e., Cobb-Douglas), they are always lower, consistent with complementarity (Rognlie 2015).

Figure 4. Housing supply elasticities and expenditure shares across regions

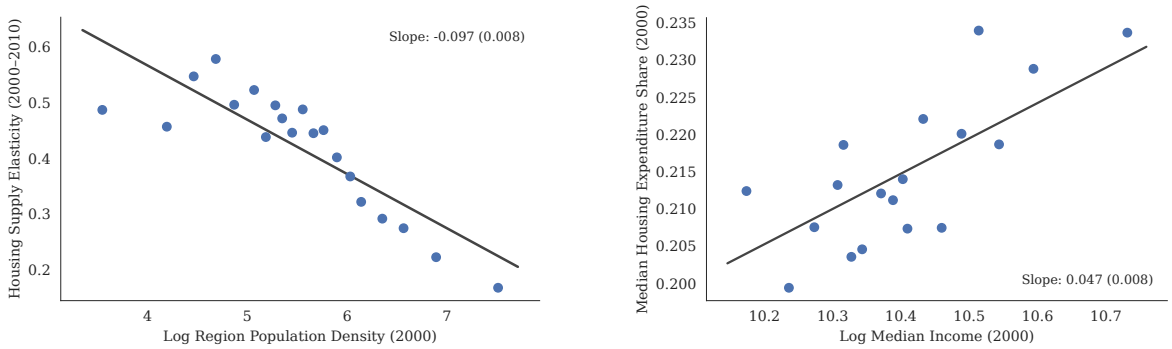

Notes: The left panel plots the relationship between Baum-Snow and Han (2024) (BSH) region-level housing-unit supply elasticities and log 2000 population density for 306 CTPP regions. Elasticities use the BSH FMM-IV specification that is quadratic in 2011 tract developed fraction and the less-tract-substitution aggregation (g_{1r} in BSH Section VI). Region population density is total Census 2000 SF1 tract population divided by tract-gazetteer land area, with both summed over tracts assigned to the region. The tract-to-region mapping is taken from BSH; because CTPP regions partially overlap, tracts assigned to multiple regions enter each assigned region. The right panel plots median housing expenditure shares against log median household wage income for 269 regions in 2000, using the MSA-level data underlying Figure 2a of Finlay and Williams (2025): 219 MSAs and 50 non-MSA state residual categories. The sample is renter households with heads aged 25–55 and at least one qualifying full-time wage worker; housing expenditure shares are annualized gross rent divided by household wage income, and both variables are region-level weighted medians. In both panels, dots are averages within equally-spaced bins of the horizontal variable. Fitted lines and reported slopes are estimated on the unbinned region-level data using OLS with heteroskedasticity-robust standard errors in parentheses.

regional variation in housing supply elasticities than regulation.[35] Second, and more importantly, the heterogeneous fixed cost model does not endogenize housing supply elasticities, which depend on exogenous distributions of fixed costs. This implies that, in response to a population decline, supply elasticities remain unchanged. In contrast, with CES housing production, supply elasticities increase everywhere as land becomes more abundant. We prefer the latter because it is conservative: (17) implies that the amount of scale dependence decreases if ζ_i^S increases in all regions.

3.2 Housing expenditure shares and PIGL demand

Next, we examine variation in housing expenditure shares across regions using estimates from Finlay and Williams (2025) and Combes et al. (2019). The former report median housing expenditure shares and median household wage income across MSAs, a coarser residential geography than the CTPP regions used above. Therefore, we examine the relationship between expenditure shares and income rather than density directly: in our model, the high-density regions that condition (ii) refers to are also the high-wage and high-price regions. The right panel of Figure 4 shows an increasing relationship, with a semi-elasticity of approximately 0.05, consistent with condition (ii) above that is needed to generate scale dependence in the direction consistent with Section 1. While

[35] Specifically, column (2) of Table 5 in Saiz (2010) implies that an interquartile change in geography, measured with the share of land unavailable for development, at average regulation, measured with the Wharton Residential Land Use Regulation Index, lowers the supply elasticity by roughly 50%. In contrast, an interquartile change in regulation at average geography lowers it by roughly 20%.

this evidence is with respect to income, Table 6 of Combes et al. (2019) reports semi-elasticities of housing expenditure shares with respect to city population in French data between 0.036 and 0.067.

The fact that housing expenditure shares vary across regions implies that we need to deviate from Cobb-Douglas demand. While a standard way of doing this would be to use CES demand, we use PIGL demand instead for three reasons.[36] First, as shown in Section 2.3, PIGL demand delivers a clean mapping between variation in expenditure shares and the price elasticity, $\zeta_i^{\mathrm{D}} = \zeta_p$:

$$\frac{p_i h_i}{w_i} = \alpha p_i^{1-\zeta_p}.$$

If p_i is positively correlated with w_i, as is true in the data and in our model, the positive slope in the right panel of Figure 4 implies that housing demand is price-inelastic, meaning $\zeta_p < 1$. Second, as discussed in Section 2.3, the key difference between CES and PIGL is whether ζ_i^{D} is constant. Given that variation in ζ_i^{D} across regions generates scale dependence, as in condition (iii) above, and we know of no evidence that ζ_i^{D} varies across regions, we prefer PIGL because it does not impose any undisciplined variation in ζ_i^{D} that would affect our main mechanism. Finally, PIGL is more tractable in our dynamic model because it delivers an explicit mapping between housing demand and prices, whereas CES requires solving for a fixed point. Nevertheless, as discussed in Section 2.3, the predictions of our model with $\zeta_p < 1$ can be understood with CES and complementarity between consumption and housing.

4 Dynamic model

In order to take the model to the data and make quantitative predictions about the effects of depopulation, we enrich the model in Section 2.3 along several dimensions. Most importantly, we make both the household and the construction-sector problems dynamic. We first describe the model environment and equilibrium, and then discuss our estimation strategy.

4.1 Model description

Environment. Table 3 lists the full set of equations summarizing the model. On the household side, we make two additions relative to the model in Section 2.3. First, each household k solves (27) subject to (29), which turns the previous static problem into a dynamic problem with migration costs and log utility, as in Caliendo et al. (2019) and Kleinman et al. (2023). The migration costs take the following form: a household that chose location i at $t-1$ faces a bilateral migration cost $\kappa_{ij} > 1$ if it moves to location $j \neq i$ at time t, with $\kappa_{ii} = 1$. Motivated by Section 3, households have

[36] Finlay and Williams (2025), Gaubert and Robert-Nicoud (2025), and Robert-Nicoud et al. (2026) incorporate PIGL preferences into spatial models. Non-homothetic PIGL is also used in models of structural transformation (Boppart 2014).

Table 3. Model environment

Block	Equation	
Households		
Value function	$V_{ikt} = \max_{j,o,c_{jokt},h_{jokt}} \{\log\left(U(c_{jokt}, h_{jokt})\epsilon_{jokt}/\kappa_{ij}\right) + \beta \mathbb{E}_t V_{jk,t+1}\}$	(27)
Taste shocks	$\epsilon_{jokt} \sim$ Nested Fréchet$(\theta, \varrho, \{z_o\})$	(28)
Budget constraint	$w_{jot} \geq c_{jokt} + p_{jt} h_{jokt}$	(29)
Indirect utility (PIGL)	$U^*(w_{jot}, p_{jt}) = w_{jot} \exp\{-\alpha(p_{jt}^{1-\zeta_p} - 1)/(1-\zeta_p)\}$ with $\zeta_p < 1$	(30)
Population dynamics	$N_t/N_{t-1} = 1 + g_t$ *(exogenous)*	(31)
	$N_{jt} = \sum_i N_{i,t-1}(1+g_t)\mathbb{P}(j_{kt} = j \mid j_{k,t-1} = i)$	(32)
Producers		
Goods production	$Y_{it} = (A_i N_{iYt}^{\sigma}) N_{iYt}$ *(external increasing returns)*	(33)
Construction firm value	$\mathcal{V}_{it} = \max_{N_{iXt} \geq 0} \{p_{it} H(X_{it}, L_i) - w_{iXt} N_{iXt} + \beta_X \mathcal{V}_{i,t+1}\}$	(34)
Structure accumulation	$X_{it} = (1-\delta) X_{i,t-1} + A_{X,i} N_{iXt}$	(35)
Housing supply	$H(X_{it}, L_i) = \left((1-\nu) X_{it}^{1-1/\eta} + \nu L_i^{1-1/\eta}\right)^{1/(1-1/\eta)}$ with $\eta < 1$	(36)
Market clearing		
Labor markets	$N_{it} = N_{iYt} + N_{iXt}$	(37)
Housing markets	$\sum_{o \in \{Y,X\}} N_{iot} h_{iot} = H(X_{it}, L_i)$	(38)
Aggregate population	$\sum_i N_{it} = N_t$	(39)

Notes: The table summarizes the equations that characterize the dynamic model.

homothetic PIGL preferences defined by the indirect utility function (30).[37]

The second addition on the household side is that households now choose both their location and their occupation, which can be goods production ($o = Y$) or construction ($o = X$). Each household receives taste shocks ϵ_{jokt} over sets of occupation-location pairs $\{o, j\}$ (e.g., goods production in Tokyo or construction in Hokkaido) drawn from a nested Fréchet distribution (28) with joint CDF

$$\mathbb{P}\left(\epsilon_{jokt} \leq \bar{\epsilon}_{jo}\ \forall j, o\right) = \exp\left\{-\sum\nolimits_j \left(\sum\nolimits_o z_o \bar{\epsilon}_{jo}^{-1/\varrho}\right)^{\varrho/\theta}\right\}, \qquad \bar{\epsilon}_{jo} > 0.$$

The parameter $\theta \geq 0$ still governs the dispersion of taste shocks across locations. The new parameter $\varrho \in [0, \theta]$ governs the dispersion between occupations, with higher ϱ (more dispersion) implying that labor supply is less responsive to wage differences. The scalars z_o shift the number of people who will choose occupation o, which will allow us to match the relatively small employment share in construction ($z_X < z_Y$).

On the production side, the goods sector produces according to (33) as before, with the change that there are increasing returns to N_{iYt}^{σ} instead of N_{it}^{σ}.[38] The main addition relative to Section 2.3

[37] (27) corresponds to a perpetual youth problem with no intergenerational linkages. Therefore, β should be interpreted as the product of the time discount factor and the survival probability. By holding β constant as we change the population, we implicitly attribute all of the population decline to changes in birth rates, which (to a first approximation) is consistent with the data (Geruso and Spears 2026).

[38] This is inconsequential, but it seems most reasonable to us that agglomeration happens within the sector.

is on the construction side, where structures are now a durable capital stock with evolution (35). Structures depreciate at rate δ and are built up with construction labor $N_{iXt} \geq 0$, which has productivity $A_{X,i}$. The durability and irreversibility of structures will generate asymmetric effects of a population decline relative to population growth (Glaeser and Gyourko 2005). The construction sector owns both the land L_i and the stock of structures X_{it}, which it combines into the CES housing stock (36) and rents out at price p_{it}, giving rise to the dynamic problem in (34).[39] As before, profits earned by developers leave the economy and are consumed elsewhere.[40]

The aggregate population level N_t is taken as exogenous and evolves according to (31). We assume that net population growth g_t happens proportionately in every location before migration decisions are made. Thus, the dynamics of population levels across locations are given by (32), where $\mathbb{P}(j_{kt} = j \mid j_{k,t-1} = i)$ are the endogenous migration rates from i to j.

To summarize, there are three substantive changes relative to Section 2.3: migration costs, which generate realistic migration rates; occupational choice, which endogenizes the elasticity of construction labor supply; and the durability of structures, which accounts for the fact that structures depreciate slowly following a population decline.

Equilibrium. We first define an equilibrium, then derive the model's equilibrium conditions.

Definition (Equilibrium)**.** *Given initial levels of population* $\{N_{j0}\}$ *and structures* $\{X_{j0}\}$ *and a time path of aggregate population growth* $\{g_t\}_{t=1}^{\infty}$*, an* ***equilibrium*** *is a time path of allocations* $\{N_{jot}, h_{jot}, c_{jot}\}$ *and prices* $\{\{w_{jot}\}, p_{jt}\}$ *that solve (27)–(39) at every* $t \geq 1$*. Given a constant population level (*$N_t = N$*,* $\forall t$*), a* ***steady state equilibrium*** *is an equilibrium in which population allocations* $\{N_{jt}\}$ *and structures* $\{X_{jt}\}$ *are constant for all* t*.*

To characterize an equilibrium, we first fix a time path of prices $\{\{w_{jot}\}, p_{jt}\}$. The dimensionality of the household's problem (27) can be dramatically reduced using the fact that the log shocks $\log \epsilon_{jokt}$ are nested Gumbel. Therefore, expected utility (i.e., before time-t shocks are realized) for a household that chose location i at time $t-1$ takes the following form, as in Caliendo et al. (2019) and Kleinman et al. (2023):

$$\widehat{V}_{i,t-1} \equiv \mathbb{E}_{t-1} V_{ikt} = \theta \log\left(\sum_j \left(\sum_o z_o w_{jot}^{\frac{1}{\varrho}}\right)^{\frac{\varrho}{\theta}} \left(\frac{v(p_{jt})}{\kappa_{ij}}\right)^{\frac{1}{\theta}} \exp\left\{\frac{\beta}{\theta}\widehat{V}_{jt}\right\}\right) + \theta\bar{\gamma}. \tag{40}$$

In (40), $v(p_{jt})$ is the second term in the PIGL indirect utility function (30) that only depends on

[39] The construction sector's discount factor β_X is not determined in equilibrium (without more assumptions) because households are hand-to-mouth. When we estimate the model, we set $\beta_X = \beta$.

[40] Letting $\Pi_{it} \equiv p_{it} H(X_{it}, L_i) - w_{iXt} N_{iXt}$ denote developer profits, aggregate goods output satisfies $Y_t \equiv \sum_i Y_{it} = \sum_{i,o} N_{iot} c_{iot} + \sum_i \Pi_{it}$.

prices, and $\bar{\gamma}$ is the Euler-Mascheroni constant. Thus, given prices, the expected value functions solve a straightforward recursion.

The same distributional properties give us labor allocations and population shares: (27) implies that, conditional on choosing location j, the share who will choose occupation o is

$$\frac{N_{jot}}{N_{jt}} = \mathbb{P}(o_{kt} = o \mid j_{kt} = j) = \frac{z_o w_{jot}^{1/\varrho}}{z_Y w_{jYt}^{1/\varrho} + z_X w_{jXt}^{1/\varrho}}. \tag{41}$$

This is a static condition because there are no costs of switching occupations over time. Wage differences can exist in equilibrium due to taste shocks $\varrho > 0$. Likewise, the outer nest of the distribution means that the share of individuals in i who move to j is

$$\mathbb{P}(j_{kt} = j \mid j_{k,t-1} = i) = \frac{\left(\sum_o z_o w_{jot}^{\frac{1}{\varrho}}\right)^{\frac{\varrho}{\theta}} \left(v(p_{jt})/\kappa_{ij}\right)^{\frac{1}{\theta}} \exp\left\{\frac{\beta}{\theta}\widehat{V}_{jt}\right\}}{\sum_{j'}\left(\sum_o z_o w_{j'ot}^{\frac{1}{\varrho}}\right)^{\frac{\varrho}{\theta}} \left(v(p_{j't})/\kappa_{ij'}\right)^{\frac{1}{\theta}} \exp\left\{\frac{\beta}{\theta}\widehat{V}_{j't}\right\}}. \tag{42}$$

In addition to lowering migration rates, the costs $\kappa_{ij} > 1$ cause migration decisions to depend on next period's expected utility $\widehat{V}_{jt}$, because households recognize that they may be "stuck" there for several periods. This means that today's location decision becomes forward-looking. Substituting (42) into (32) and dividing by N_t gives the evolution of population shares

$$\frac{N_{jt}}{N_t} = \sum_i \frac{N_{i,t-1}}{N_{t-1}} \underbrace{\frac{\left(\sum_o z_o w_{jot}^{\frac{1}{\varrho}}\right)^{\frac{\varrho}{\theta}} \left(v(p_{jt})/\kappa_{ij}\right)^{\frac{1}{\theta}} \exp\left\{\frac{\beta}{\theta}\widehat{V}_{jt}\right\}}{\sum_{j'}\left(\sum_o z_o w_{j'ot}^{\frac{1}{\varrho}}\right)^{\frac{\varrho}{\theta}} \left(v(p_{j't})/\kappa_{ij'}\right)^{\frac{1}{\theta}} \exp\left\{\frac{\beta}{\theta}\widehat{V}_{j't}\right\}}}_{\mathbb{P}(j_{kt}=j|j_{k,t-1}=i)}. \tag{43}$$

Without moving costs ($\kappa_{ij} = 1$), origin dependence drops out of the migration probabilities, and in a steady state this collapses to the static shares in Section 2.3.

The conditions (40)–(43) determine allocations given prices. To determine prices, we use the firms' optimization and market clearing. The fact that the goods sector earns zero profits implies

$$w_{iYt} = A_i N_{iYt}^{\sigma}.$$

The construction sector's first-order and envelope conditions of (34) give

$$\frac{w_{iXt}}{A_{X,i}} = p_{it}\frac{\partial H(X_{it}, L_i)}{\partial X_{it}} + \beta_X(1-\delta)\frac{w_{iX,t+1}}{A_{X,i}}, \tag{44}$$

which is similar to the static condition, except we now have discounting and depreciation.[41] Finally,

[41] In our solution algorithm described in Appendix D, we account explicitly for the non-negativity constraint $\hat{N}_{iXt} \geq 0$.

to determine house prices, we take the households' housing demand curves $h_{iot} = \alpha w_{iot} p_{it}^{-\zeta_p}$, multiply both sides by N_{iot}, sum over o, and use (38):

$$p_{it} = \left(\alpha \frac{\bar{w}_{it}}{H(X_{it}, L_i)/N_{it}} \right)^{1/\zeta_p}, \quad \bar{w}_{it} \equiv \sum_o \frac{N_{iot}}{N_{it}} w_{iot}.$$

Naturally, prices are higher in locations with higher average income $\bar{w}_{it}$ and lower per-capita housing supply, and these differences are amplified when $\zeta_p < 1$.

Solution method. We solve for the steady state equilibrium numerically using outer iteration on the equilibrium condition for population shares in (43), with inner iterations on (40) to solve for households' value functions and (44) to solve for the employment share of the construction sector. Given initial and terminal steady states that are defined by two different aggregate population levels, we then solve for the equilibrium transition path with perfect foresight using a forward-backward algorithm. Appendix D describes both algorithms in detail.

4.2 Model estimation

External parameters. We set the model period to one year, the set of regions to be the Japanese prefectures described in Section 1.1 (as in Giannone et al. 2026), and calibrate five parameters externally based on existing estimates in the literature. First, we follow Kleinman et al. (2023) and set the household and landlord discount factors to $\beta = \beta_X = 0.95$.[42] Second, we use an inverse migration elasticity of $\theta = 2.02$ based on the implied annual elasticity from the estimation in Caliendo et al. (2019), which is close to the estimate of 2.08 from Greaney et al. (2025). Third, we set the agglomeration elasticity to $\sigma = 0.05$, which is the mean of the estimates surveyed by Ahlfeldt and Pietrostefani (2019) for the elasticities of labor and total factor productivity to density.[43] Finally, we set the price elasticity of housing demand to $\zeta_p = 0.75$, which is the midpoint of the range of estimates reported in Robert-Nicoud et al. (2026) that come from variation in housing expenditure shares across regions.

Estimation procedure. We estimate the remaining parameters to match prefecture population shares and the moments in Table 4. We compute the model's moments in a steady state with the aggregate population level equal to Japan's population in 2000. We normalize construction

However, this constraint does not bind in equilibrium because extreme value taste shocks imply that a positive mass of households will always work for any positive wage.

[42] In reality, β_X and β could differ because of mortality and the fact that demographic changes lead to changes in interest rates (Auclert et al. 2021). However, β_X plays a minimal role in our experiments because it does not affect relative prices, so we prefer to follow Kleinman et al. (2023) and keep it equal to β.

[43] The value of σ is sufficiently small that it does not play an important role in our quantitative results. All of our results are quantitatively similar if we instead use $\sigma = 0$.

productivity $A_{X,i} = 1$ across regions; then, to put land into the same units as structures, we take the inhabitable land area $\{L_i\}$ from prefecture-level data and scale them so that $\sum_i L_i = N_0$.[44] Our estimation procedure then iterates between two steps until convergence, which are described in detail in Appendix E.

Given a guess of the parameters, the first step is to invert the productivities $\{A_i\}$ from the prefecture population shares in 2000. Given the inverted $\{A_i\}$, the second step is to choose the remaining parameters so that the model exactly matches the set of moments in Table 4. Four moments are specific to Japan: (i) the (population-weighted) average prefecture out-migration rate over 1990–2000 to identify the bilateral moving cost, κ_{ij}, which we assume is constant and equal to κ for all $i \neq j$; (ii) the aggregate Japanese housing depreciation rate over 1995–2024 to identify the structures depreciation rate, δ; (iii) the Japanese construction-sector factor income share to identify the land share in housing production, ν; and (iv) the ratio of the construction-sector employment to factor income shares to identify the occupation preference shifters, z_Y/z_X.[45] We use the housing supply elasticities from Baum-Snow and Han (2024) (BSH) that were discussed in Section 3. In particular, we identify ϱ by matching the housing supply elasticity in Tokyo in the model to the 10th percentile of the distribution of estimates from BSH in Figure 4; see (E.11) for an expression for the model's supply elasticity as a function of ϱ.[46] This is conservative, because housing supply elasticities decrease with density, and Tokyo is denser than any region in the BSH sample.[47] We then identify η from the slope coefficient in the left panel of Figure 4, as illustrated in (26). Third, we identify α by targeting the aggregate housing expenditure share, which we take from Piazzesi and Schneider (2016). The final moment is a normalization: because preferences are homothetic, aggregate GDP per capita can be normalized to any value by changing the scale of A_i. We choose a value of one, so that all variables are relative to aggregate GDP per capita in 2000.

Estimation results. Table 4 shows our parameter estimates and the comparison of model and data moments. Our model is able to replicate the data moments almost exactly, given the model is just-identified. Importantly, we estimate that a value of $\eta = 0.65$ is required to match the relationship

[44] These two normalizations of scale are without loss because housing production is homogeneous of degree one in (X_{it}, L_i), so rescaling the units of land by a common factor c can be offset by measuring structures in the same units, $X'_{it} = cX_{it}$, and setting $A'_{X,i} = cA_{X,i}$. This transformation leaves X_{it}/L_i, the land and structure output shares, and the housing-supply elasticities unchanged. It also leaves the allocation unchanged after expressing housing prices as $p'_{it} = p_{it}/c$, and rescaling $\alpha' = \alpha c^{1-\zeta_p}$ so that expenditures are unchanged. With homothetic PIGL preferences, the latter change in indirect utility is common across locations and occupations and therefore does not affect choices.

[45] (41) shows that we can normalize $z_X = 1$ without loss of generality.

[46] We compute the model's supply elasticity in response to a demand shock that affects prices without affecting migration (i.e., the partial equilibrium slope of the housing supply curve). This is the consistent with the estimates in BSH, which are identified using differences in housing demand within regions.

[47] By "conservative," we mean that this choice reduces the amount of scale dependence. This is because (17) implies that differences in housing supply elasticities generate larger differences in population shares when housing supply elasticities are lower.

Table 4. Parameter estimates and model fit

Parameter	Value	Moment	Source	Model	Data
κ	$10^{6.45}$	Average out-migration rate	Japan 1990-2000	0.024	0.024
ϱ	1.20	Tokyo housing supply elasticity	Baum-Snow and Han (2024)	0.210	0.210
δ	0.08	Housing depreciation rate	Japan 1995-2024	0.058	0.058
ν	0.26	Construction factor income share	Japan 2000	0.085	0.085
z_Y/z_X	9.23	Construction empl. / income	Japan 2000	1.150	1.151
η	0.65	Housing supply vs. density slope	Baum-Snow and Han (2024)	-0.097	-0.097
α	0.34	Housing expenditure share	Piazzesi and Schneider (2016)	0.250	0.250
$\sum_i A_i$	16.67	Aggregate GDP per capita	Normalization	1.002	1.000

Notes: The table reports the parameter estimates from the baseline estimation, as well as the fit of the model on the targeted moments described in the main text. See Appendix E for a description of how each moment is computed in the model.

between housing supply elasticities and population density in the left panel of Figure 4. While that relationship was targeted, Figure F.1 shows that the model also provides a good fit to the relationship between land share of housing value and region density discussed in Section 3. The inverted values of $\{A_i\}$ are reported in Figure F.2; Figure F.3 shows we match the 2000 population shares well. The highest values of A_i are estimated in Tokyo and Osaka, as well as their surrounding areas, given that these prefectures have high population shares, despite having small land areas (Figure F.4).

Alternative models. We also estimate three alternative models that correspond to alternative assumptions about housing supply and demand: (i) Cobb-Douglas housing supply with Cobb-Douglas housing demand, in which we impose $\eta = \zeta_p = 1$; (ii) CES housing supply with Cobb-Douglas housing demand, in which we impose $\zeta_p = 1$ but estimate η; and (iii) Cobb-Douglas housing supply with PIGL housing demand, in which we impose $\eta = 1$ but keep $\zeta_p = 0.75$. The fit of each alternative model on the targeted moments is reported in Table F.1, Table F.2, and Table F.3, respectively. We use these alternative estimated models in the next section to isolate the role of housing supply versus demand in shaping the response to population decline.

5 Quantitative results: Depopulation in Japan

In this section, we use our estimated model to study the quantitative consequences of Japan's depopulation. We first compare two steady states that correspond to different aggregate population levels, and then study the transition between them.

Aggregate population path. When we estimate the model in Section 4.2, we use Japan's aggregate population level and prefecture-level population shares in 2000, which we treat as the initial steady state. To determine the transition path and terminal steady state, we use the medium-fertility / medium-mortality population projection from National Institute of Population and Social Security

Figure 5. Transition path of aggregate Japan population

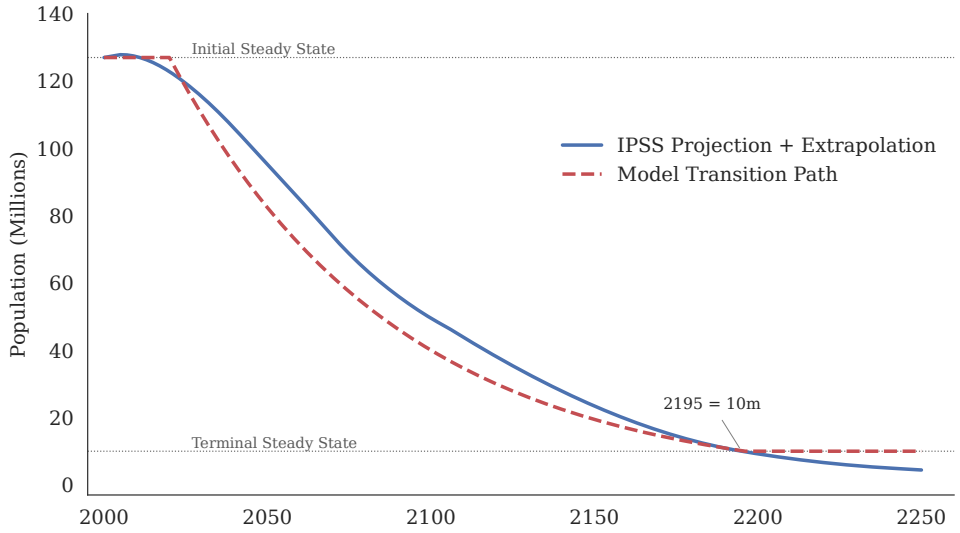

Notes: The figure plots two aggregate population paths. The solid blue line is the medium-fertility / medium-mortality projection of National Institute of Population and Social Security Research (2006). We construct this series using the 2005–2055 main projection in Table 1-1 and the 2056–2105 conditional continuation in Table 3-1. To extrapolate outside of this window, we combine these series with Japan's 2000 census population prior to 2005 and a monotone cubic interpolation after 2105 through the 2195, 2240, and 2350 milestones reported in Gietel-Basten (2016). The conditional continuation extends the same medium-fertility / medium-mortality variant under the additional assumption that fertility returns to replacement level by 2155. The dashed red line is the path that we use in our model transition, which is described in the main text.

Research (2006) (IPSS). As of 2005, this forecast projected that Japan's population would shrink to slightly below its 2000 level by 2020, close to the realized decline that matched the 2000 level. Motivated by this, we consider a transition path in which the aggregate population remains constant at its 2000 level for the first 20 years, then declines at a constant rate until 2195, and is constant at the terminal level thereafter. We set the rate of decline such that the aggregate population reaches 10 million by 2195, the level implied by the extrapolation of the IPSS projection in Gietel-Basten (2016).

Figure 5 compares our transition path to the IPSS projection and extrapolation. In principle, we could use the IPSS projection directly, but we use the path described above for three reasons. First, an initial period in which the population is flat allows us to cleanly study how households respond in anticipation of the future decline, without confounding from contemporaneous changes in the aggregate population. Second, truncating the decline at a finite terminal population allows us to study how long it takes the model to converge to the terminal steady state after the decline stops.[48] Third, a constant rate of decline avoids the non-monotone growth-rate profile of the IPSS projection.

[48] We could instead consider a terminal steady state with zero population, which is what the IPSS projection implies. However, reaching this steady state takes a very long time, and the computational cost of solving for the transition path increases with its length.

Figure 6. Change in prefecture population shares across steady states

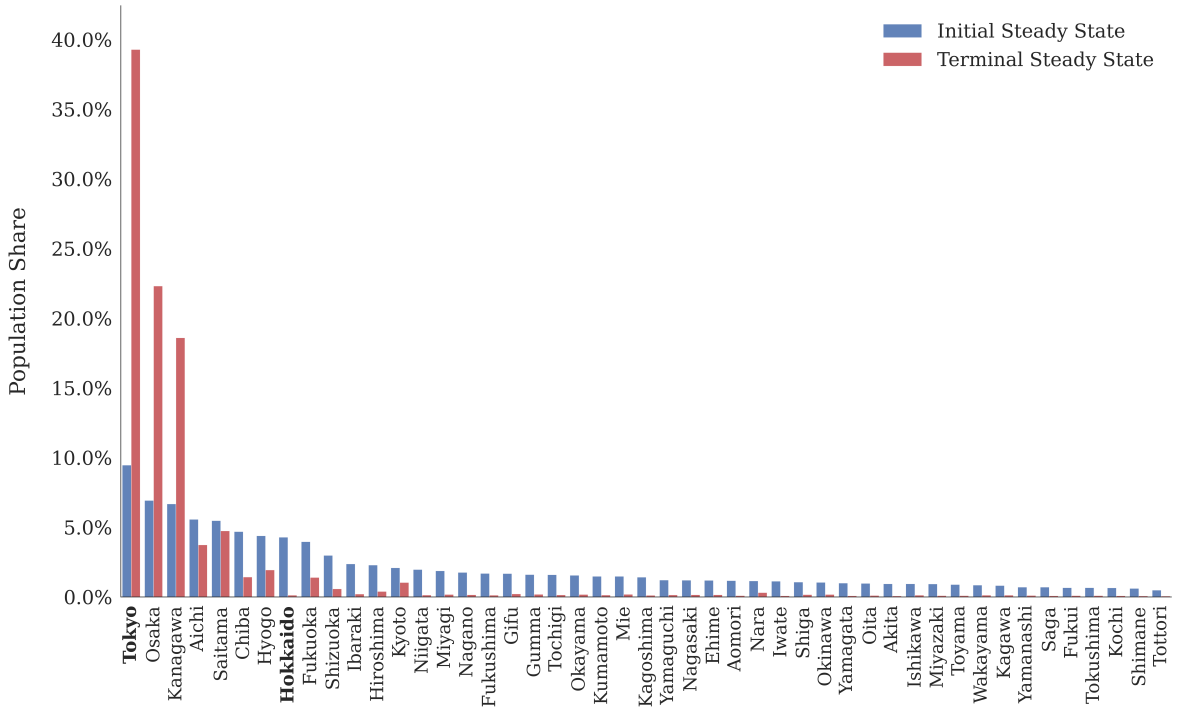

Notes: The figure reports each prefecture's population share from the baseline model in the initial steady state (blue) and the terminal steady state (red). Prefectures are sorted by their initial population share.

5.1 Comparison across steady states

Population shares. Figure 6 shows the population shares of each prefecture in the initial and terminal steady states, with prefectures sorted by their initial population share. In response to population decline, Tokyo's share quadruples from approximately 10% to nearly 40%. In terms of levels, Tokyo's population falls from 12 million to 4 million, relative to approximately 1 million with scale independence. While the population shares of Osaka and Kanagawa also increase and more than double, the remaining prefectures experience significant declines, with many emptying out almost completely. An extreme case is Hokkaido, whose share falls from over 4% to 0.1%.

The substantial heterogeneity in the changes in population shares across prefectures is determined by initial population density, as in Section 2.3. Panel A of Figure 7 plots the log change in each prefecture's population share against its initial density measured by its log population share scaled by inhabitable land share, which is the comparative static in (23). The green points show the results in the baseline model, in which this relationship is monotonic and approximately linear: the densest prefectures gain the most and the sparsest lose the most. The two shaded circles on the left and the right are Hokkaido and Tokyo, which are the least and most dense regions, respectively.

Figure 7. Changes across steady states by density

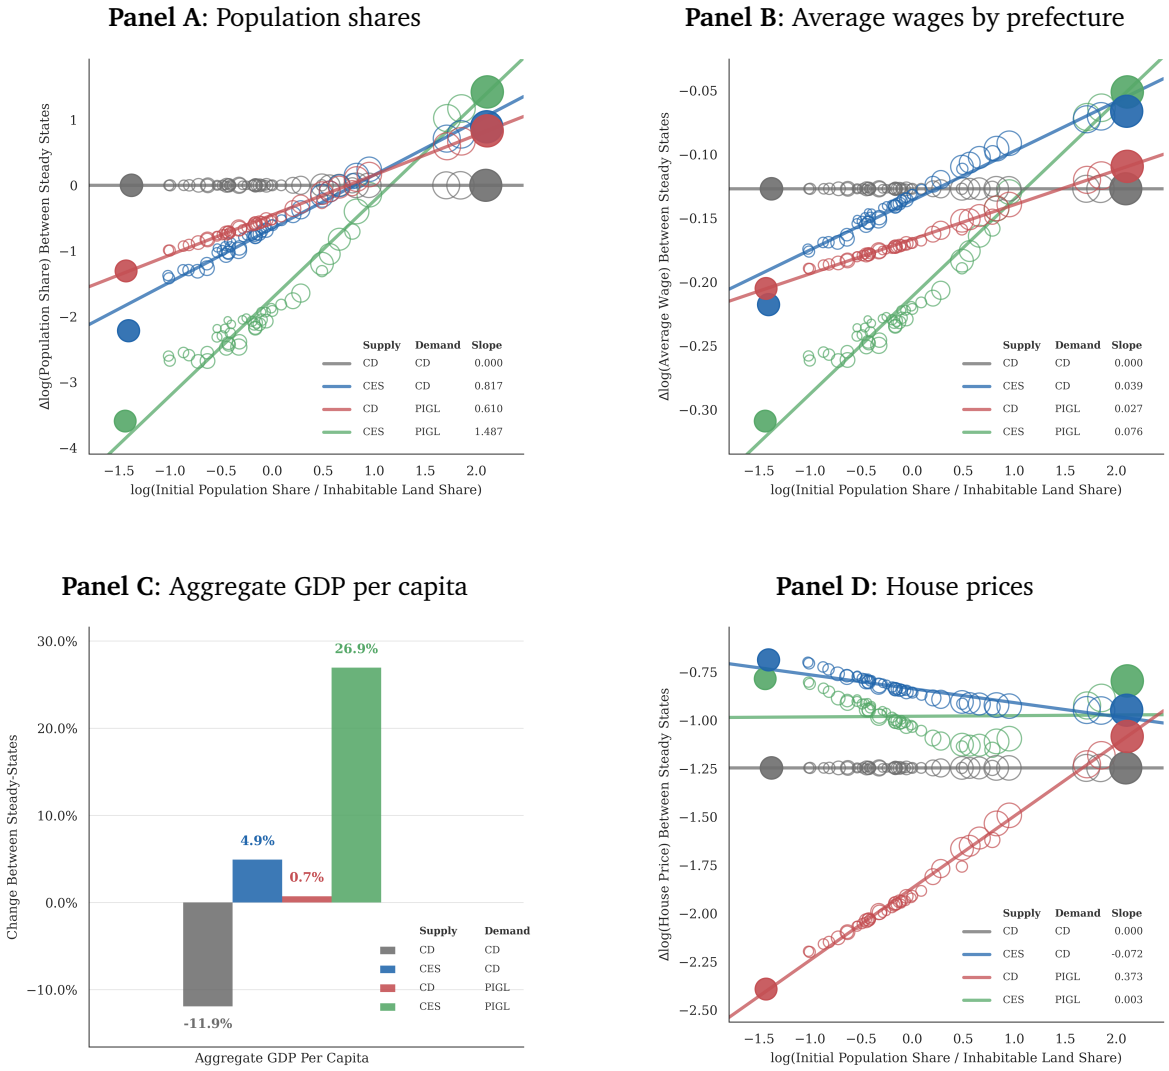

Notes: Panels A, B, and D plot percent changes between the initial and terminal steady states of different variables against the prefecture's initial steady state log population share scaled by inhabitable land share, $\log((N_i/N)/(L_i/\sum_j L_j))$. Each marker represents one prefecture, where the rightmost shaded circle corresponds to Tokyo and the leftmost shaded circle corresponds to Hokkaido. Panel A plots $\Delta \log(N_i/N)$, Panel B plots $\Delta \log \bar{w}_i$ where $\bar{w}_i$ is the prefecture's population-weighted average wage across the goods and construction sectors, and Panel D plots $\Delta \log p_i$. Marker areas are proportional to initial steady state population, and fitted lines are population-weighted regression lines with the slopes reported in the legend. Panel C reports the percent change in aggregate GDP per capita between the initial and terminal steady states. Across all panels, colors index the four estimated models indicated in the bottom right of each figure.

The remaining three sets of points in Panel A repeat the same exercise in the three alternative models estimated in Section 4.2, all of which have the same initial population distribution by construction. The gray points correspond to the model with Cobb-Douglas housing production (i.e., supply) and preferences (i.e., demand). As expected, this model features scale independence, with no change in population shares across steady states. The blue points show that introducing CES

housing supply generates the same pattern as the baseline model, with populations falling by less in denser regions. However, quantitatively, the slope of this relationship is around $0.817/1.487 = 55\%$ of the baseline model. Finally, the red points show that using PIGL demand but Cobb-Douglas supply generates a similar relationship with a slope of around $0.610/1.487 = 41\%$ of the baseline model. Therefore, CES supply contributes slightly more to scale dependence than PIGL demand, and the interaction between the two is small.

Average wages. Panel B of Figure 7 shows that the log change in average wages across prefectures follows a similar pattern to the population shares in Panel A. Average wages in the baseline model fall everywhere because the loss of agglomeration is global, but they fall by much less—around 5%—in Tokyo relative to close to 25% in Hokkaido. To understand these responses, recall (24) from our static model: the change in wages is a scaled version of the change in population shares, where the scale is $d \log \bar{w}_i / d \log N_i$. Given the relatively small and stable construction employment share, this term is approximately the elasticity of the goods wage to the local population, σ. This is exactly what the slope of this relationship shows: dividing the wages slope in the baseline model by the population-shares slope gives $0.076/1.487 \approx 0.05 = \sigma$.

Aggregate output. Despite the fact that wages fall everywhere, Panel C shows that aggregate GDP per capita rises by 27% in the baseline model, which mostly comes from the goods sector (Figure F.5). To understand the source of this increase, the goods-output component of GDP per capita can be written as

$$\frac{Y_t}{N_t} = \underbrace{\frac{N_{Yt}}{N_t}}_{\text{goods emp. share}} \times \underbrace{N_{Yt}^{\sigma}}_{\text{scale}} \times \underbrace{\sum\nolimits_i A_i \left(\frac{N_{iYt}}{N_{Yt}} \right)^{1+\sigma}}_{\text{allocation}}$$

where $N_{Yt} = \sum_i N_{iYt}$. The goods employment share converts output per goods-sector worker into a contribution to aggregate GDP per capita. The scale term captures the loss of agglomeration as the aggregate population falls. The allocation term reflects the sorting of households across prefectures: it is a productivity-weighted average that puts more mass on prefectures with higher goods-sector population shares. With scale independence, the allocation term is unchanged because population shares do not change, so the scale term dominates and causes aggregate GDP per capita to fall by approximately 12%. In the baseline model with scale dependence, households reallocate toward Tokyo and Osaka, which have the highest values of A_i, and the rise in the allocation term more than offsets the scale term. With just CES supply or PIGL demand, GDP per capita still rises but by less.

House prices. Panel D of Figure 7 shows the log change in house prices across prefectures.[49] Unlike population shares and wages, which fall by less in high-density prefectures, house prices fall by similar amounts across prefectures in the baseline model. To understand why the fall in house prices is similar even with scale dependence, recall (25), which shows that there are two offsetting forces that affect house prices. The first is the effect of CES supply, which creates heterogeneity in the local price response $d \log p_i / d \log N_i$: prices fall more in dense prefectures because housing supply per capita goes up by more. Because this shows up in both the numerator and the denominator, this implies that prices should fall by more in dense prefectures, holding $p_i h_i / w_i$ constant. The blue points in Panel D show this is the case in the model with CES supply but Cobb-Douglas demand.

However, (25) shows that there is a second force from the demand side that works in the opposite direction. PIGL demand creates heterogeneity in the housing expenditure share, with higher shares in dense prefectures. Given the expenditure share shows up in the denominator, this implies that, holding $d \log p_i / d \log N_i$ constant, prices should fall by less in dense prefectures. This is consistent with the red points in Panel D, which show results for the model with Cobb-Douglas supply and PIGL demand. Intuitively, Cobb-Douglas supply means that housing supply per capita increases everywhere by the same amount, holding migration choices fixed. PIGL demand then creates an equilibrium migration response that drives up the population shares of dense prefectures, causing their housing supply per capita and, hence, prices to fall by less.

5.2 Transition between steady states

We now turn to the transition path between the initial and terminal steady states. We consider a perfect foresight transition in which the aggregate population path in Figure 5 is announced in 2000, with the decline starting in 2020. We focus on two time periods along the transition: the anticipation window (2000–2020), during which the aggregate population is constant, and the long run, during which the economy converges to the terminal steady state.

Anticipation. Panels A and B of Figure 8 show the cumulative changes in population shares and house prices over 2000–2020. Tokyo's population share rises by approximately 3% and Hokkaido's falls by approximately 3%. We see similar patterns in house prices, with prices increasing by 3% in Tokyo and falling by 1% in Hokkaido. These anticipation effects prior to the aggregate population decline reflect two forces. First, households face migration frictions and recognize that they may be "stuck" in their location for several periods, so the option value of moving to Tokyo today rises as soon as the future population decline is announced. Second, the construction sector is forward-looking and slows investment everywhere given the population decline, but slows it by more in prefectures

[49] Figure F.6 shows the changes in the remaining endogenous variables across steady states.

Figure 8. Population shares and house prices during the transition

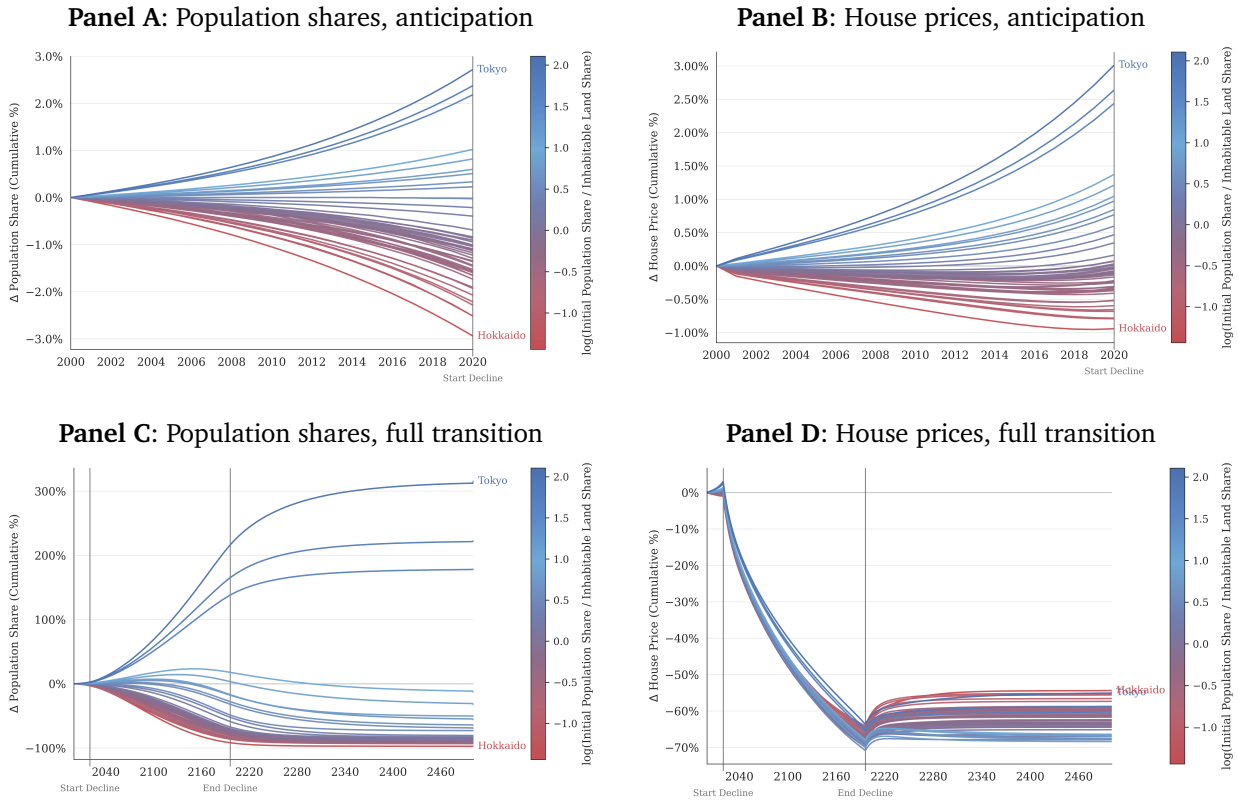

Notes: The figure plots the transition path of the baseline model. Each line corresponds to one prefecture, color-coded by initial log population share over inhabitable land share, $\log((N_i/N)/(L_i/\sum_j L_j))$. Panels A and B cover 2000–2020, the period before the aggregate population begins to decline; panels C and D cover the full transition through the terminal steady state. The vertical axis is the cumulative percent change from the initial steady state. Vertical markers indicate the start and end of the aggregate population decline shown in Figure 5.

where it foresees future depopulation. This decrease in housing supply coupled with the increase in migration drives up house prices in Tokyo. In Hokkaido, the out-migration drives down prices, but the decrease in housing supply means prices fall by less than they rise in Tokyo.

Comparison with data. Given the actual population in Japan was unchanged between 2000 and 2020, we can compare the empirical relationship between the change in population shares and initial density from Figure 2 to what happens in our model. Figure 9 reproduces Figure 2 pooling 2000–2019 in the model (left) and in the data (right).[50] The results show that the model generates a similar qualitative pattern to the data, where denser prefectures experienced larger increases in population shares. Quantitatively, the model generates a slope of approximately 0.017 relative to a data slope of approximately 0.079, explaining approximately one-quarter of the observed rise in spatial concentration. We view this as a meaningful share of the observed reallocation, given that the

[50] We choose 2019 instead of 2020 to avoid any effects of COVID-19.

Figure 9. Initial population density and subsequent population growth: model versus data

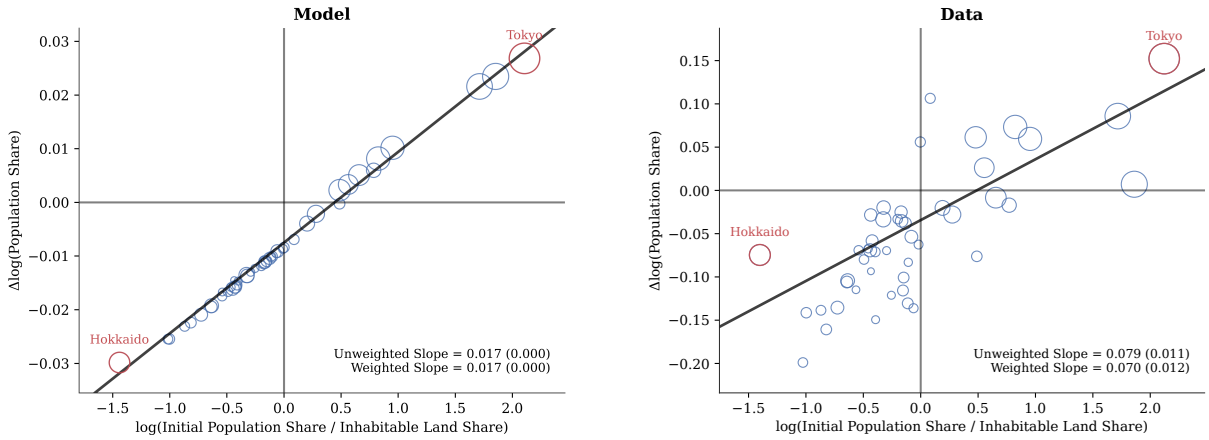

Notes: The figure compares the relationship between initial population density and subsequent population-share growth in the model and the data. Each marker represents one prefecture. The horizontal axis is the log of initial population share divided by inhabitable land share, $\log((N_i/N)/(L_i/\sum_j L_j))$, and the vertical axis is the change in log population share from 2000 to 2019. The left panel uses the baseline model transition path, while the right panel uses prefecture population data from the Statistics Bureau of Japan. Marker areas are proportional to 2000 population. The black fitted line is the initial-population-weighted regression line; the legend reports unweighted and initial-population-weighted slopes with heteroskedasticity-robust standard errors in parentheses.

model contains only two mechanisms related to housing. Importantly, this exercise holds prefecture fundamentals constant and imposes constant birth and death rates across prefectures, while some of Tokyo's population increase has reflected differential birth rates (Figure A.2, Giannone et al. 2026). In Section 5.3, we discuss some other mechanisms that may be at work in the data.[51]

Full transition. Panels C and D of Figure 8 show the same objects as Panels A and B over the full transition between the two steady states.[52] Population shares continue to diverge during the population decline, but take over 200 years after the end of the decline to reach the new steady state. This long lag reflects the slow depreciation and irreversibility of structures. House prices, while rising in anticipation in Tokyo, begin falling everywhere once the population decline begins because the increase in housing supply per capita is so large. Additionally, once housing supply per capita starts increasing, the relationship between house price changes and initial density flattens out, as described in (25).

5.3 Extensions

We briefly discuss extensions to the baseline model, with additional details in Appendix F.

[51] One mechanism that is missing from our model is a notion of distance between prefectures. However, Figure F.9 shows that the model-implied and observed origins of net migration to Tokyo over 2010–2019 are not extremely different.

[52] Additional variables along the transition path are reported in Figure F.7 and Figure F.8.

Spending commitments. Another reason why sparse areas may shrink in response to depopulation is because local public goods, such as parks, schools, or firefighters, become harder to finance. To capture this idea, we extend our model to include public goods that are financed with proportional wage taxes in each prefecture. Local governments all choose the same constant steady state tax rate, but we assume that spending adjusts slowly toward this target. This sluggish response of spending has no effect on the change in population shares across steady states because spending eventually adjusts in the long run. However, along the transition path, tax rates rise most in less dense prefectures (e.g., Hokkaido), which have fewer residents to finance previously committed spending. These differences in tax rates speed up the transition, amplifying the population share and house price responses prior to the actual population decline by around 15%.

Amenity spillovers. While our model features spillovers in production, existing literature has also emphasized spillovers in amenities (e.g., Diamond 2016; Fajgelbaum and Gaubert 2020; Couture et al. 2024). We incorporate these spillovers by introducing a reduced-form term in flow utility with a constant elasticity with respect to local population, as in Ahlfeldt et al. (2015) and Giannone et al. (2026). Using a negative value of this elasticity that is around the median of existing estimates dampens the reallocation toward dense prefectures, but does not overturn it: Tokyo's population share still doubles in response to the population decline. In contrast, a positive amenity spillover, which Giannone et al. (2026) provides evidence for in Japan, would strengthen our baseline results.

Endogenous fertility. We take the population decline as given, but fertility may respond endogenously to depopulation: as housing supply expands, households may have more children, and the aggregate population may stabilize at a positive level. In Appendix F.5, we consider what happens if birth rates depend on per capita housing consumption, consistent with empirical evidence (Lovenheim and Mumford 2013; Dettling and Kearney 2014; Fazio et al. 2025). Two qualitative results emerge. First, the decline in the aggregate population is indeed dampened, but the dampening is smaller in dense prefectures where housing supply expands less. Second, when housing production is CES with $\eta < 1$, per-capita housing supply is bounded above: as the population falls, structures become the "weak link" and housing supply reaches an upper bound. As a result, the birth rate may never rise enough to stabilize the aggregate population, which can eventually fall to zero. In contrast, with the Cobb-Douglas housing production, per capita housing supply increases without bound as the population falls, which creates a stable positive population level.

6 Optimal spatial policy response to depopulation

6.1 Planner's problem

We consider a social planner who maximizes the average utility of all households by choosing transfers across locations and occupations that are budget-neutral. For a given aggregate population level, N, the planner's problem in a steady state equilibrium (dropping t subscripts) is

$$W^{\text{SP}}(N) \equiv \max_{\{\tau_{io}\}} \sum_i \frac{N_i}{N} \widehat{V}_i \tag{45}$$
$$\text{subject to:} \qquad 0 \geq \sum_{i,o} N_{io} \tau_{io},$$

where N_i and $\widehat{V}_i$ implicitly depend on the entire vector of transfers. These transfers must be incentive-compatible, so the planner can only manipulate allocations indirectly through them. We remove any inefficiency coming from profits in the construction sector by assuming that they are redistributed evenly to all households (Redding and Rossi-Hansberg 2017). Therefore, letting $\bar{\pi} \equiv \sum_j \Pi_j / N$ denote national per-capita profits, households' budget constraints become

$$w_{jo} + \bar{\pi} + \tau_{jo} \geq c_{jo} + p_j h_{jo}.$$

In our dynamic model, we consider a Ramsey-style problem in which we restrict the planner to choosing a wage-indexed transfer rule of the following form:

$$\tau_{io} = \tau_w w_{io} + \tau_0. \tag{46}$$

This policy rule has two parameters: a lump-sum transfer (or tax), τ_0, and a constant transfer rate on wages, τ_w. A positive (negative) value of τ_w corresponds to making transfers to high-wage (low-wage) locations, which generally have larger population shares. We choose (46) because, as shown below, the optimal location-specific transfers in our static model take this form.

6.2 Optimal policy results

Optimal transfers. Figure 10 shows the transfers that solve (45) in our dynamic model. Panel A plots the average transfer against population density across prefectures in the initial steady state. Across all models, the planner makes transfers from sparse prefectures (e.g., Hokkaido) to dense prefectures (e.g., Tokyo). Quantitatively, the optimal τ_w is similar across the models and around 20%, which imply transfers of around 10% of aggregate GDP in Tokyo and -3% in Hokkaido. In

Figure 10. Optimal steady state transfers

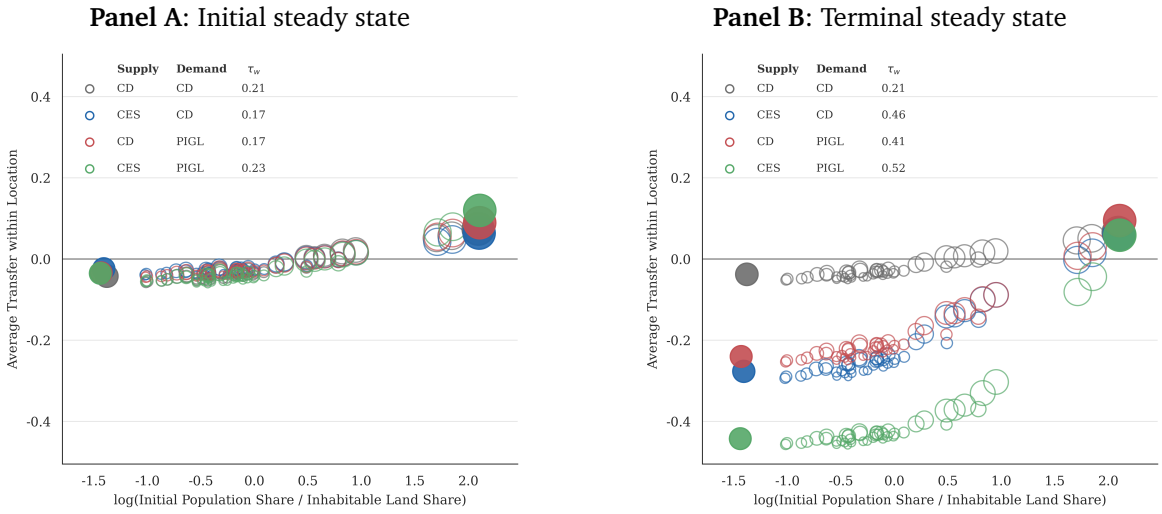

Notes: The figure reports the planner's optimal steady state wage-indexed transfer policies across the four estimated model specifications. Panels A and B plot prefecture-average transfers net of the common profit rebate, $\tau_w(\bar{w}_i - \bar{w})$, in the initial and terminal steady states, respectively, where $\bar{w}_i$ is the average wage in prefecture i and $\bar{w}$ is the national average wage. The horizontal axis is the prefecture's initial steady state log population share scaled by inhabitable land share, $\log((N_i/N)/(L_i/\sum_j L_j))$. Each marker represents one prefecture; marker areas are proportional to initial steady state population, and the shaded circles correspond to Tokyo and Hokkaido. The embedded legends report the housing-supply and housing-demand specifications and the optimal value of τ_w in (46). Colors index the four estimated models indicated in the legends.

contrast, Panel B shows that in the terminal steady state, the optimal transfers differ dramatically across models. With scale independence (Cobb-Douglas housing supply and demand), optimal transfers do not change relative to the initial steady state, given that equilibrium population shares remain unchanged. With scale dependence, τ_w increases by a factor of more than two, reaching 52% in the baseline model, with transfers of approximately 5% of aggregate GDP in Tokyo and -45% in Hokkaido.

Welfare gains. The welfare gains from the optimal transfers are shown in Figure F.10. We measure the welfare gains in consumption-equivalent units relative to the equilibrium without transfers (but with profit redistribution). In the initial steady state, the welfare gains are small in all four models at 0.2–0.3% of lifetime consumption. However, in the terminal steady state, the welfare gains become an order of magnitude larger in the models with scale dependence. In the baseline model (CES and PIGL), the gains are equivalent to an increase of 5.1% of lifetime consumption.

Mechanisms. To understand these results, we first characterize optimal policy in our static model from Section 2.3 and then compare it with our dynamic model. Assuming that the planner can only make transfers across locations, the optimal transfers in the static model take the following form.

Proposition 5 (Optimal policy in the static model, Fajgelbaum and Gaubert 2020)**.** *In the general*

static model (Section 2.3), the planner's optimal per-capita transfer to location i is

$$\tau_i = \underbrace{-\frac{\theta}{\theta+1}(w_i - \bar{w})}_{\text{redistribution}} + \underbrace{\frac{\sigma}{\theta+1}\left(\frac{Y_i}{N_i} - \frac{Y}{N}\right)}_{\text{agglomeration}}, \tag{47}$$

where $\bar{w} \equiv \sum_j N_j w_j / N$ is average income and $Y/N = \sum_j Y_j/N$ is per-capita goods output.

The optimal transfers (47) are the same as those derived by Fajgelbaum and Gaubert (2020) and balance two competing forces.[53] On the one hand, the planner wants to tax high-income places (Tokyo) and subsidize low-income places (Hokkaido) as a form of insurance against taste shocks. On the other hand, because goods producers do not internalize the agglomeration benefits $\sigma > 0$, there is an inefficiently low number of workers in high-output areas, so the planner wants to subsidize Tokyo to attract workers there. Importantly, although housing creates congestion, it does not show up in (47) because it is priced by the competitive housing sector that returns profits to households.

Given the values of θ and σ in our dynamic model, Proposition 5 implies that redistribution should dominate agglomeration, resulting in $\tau_w < 0$. However, as shown in Figure 10, solving (45) in our dynamic model delivers a positive τ_w. In Appendix F.6, we progressively remove the differences between our dynamic model and the static model in which Proposition 5 is derived: moving costs, depreciation, and occupation-specific taste shocks. We find the change in sign of τ_w is driven by moving costs. To see why moving costs affect optimal transfers, differentiate the planner's objective in (45) to obtain:

$$\frac{dW^{\text{SP}}}{d\tau_w} = \sum_i \frac{N_i}{N}\frac{d\widehat{V}_i}{d\tau_w} + \sum_i \frac{dN_i/N}{d\tau_w}\widehat{V}_i$$

This equation shows that the planner can affect welfare in two ways: changing the utility of people living in different locations (the first term) and moving people to locations that have higher average utility (the second term). Without moving costs, households' location decisions are static, so $\widehat{V}_i$ is constant across i and the second term is zero. With moving costs, $\widehat{V}_i$ varies and the planner can improve welfare in steady state by reallocating the population from low- to high-utility places.

These differences in utility across prefectures are the key reason that the optimal τ_w is positive in the initial steady state and increases in the terminal steady state. Figure F.11 shows that the utility of living in dense prefectures (e.g., Tokyo) is higher than in sparse prefectures (e.g., Hokkaido) in the initial decentralized steady state, creating gains from moving people to the former. These

[53] There are two minor differences from Fajgelbaum and Gaubert (2020). First, the agglomeration externality in (47) only applies to goods output Y_i/N_i, not income to all sectors w_i, because we have assumed that the construction sector does not enjoy the agglomeration externality. Second, Fajgelbaum and Gaubert (2020) also show that negative amenity externalities—non-priced spillovers like crime and congestion—generate large additional externalities that cause the planner to want to tax cities and subsidize sparse areas. When we prove Proposition 5 in Appendix B.5, we allow for an exogenous spillover $N_i^{-\gamma}$ and show that (47) is the same, except θ is replaced with $\theta + \gamma$.

differences are initially similar across the different models because all models are estimated to match initial population shares, which explains why their policies are similar. However, scale dependence amplifies these utility differences in the terminal steady state because housing becomes disproportionately less scarce in dense prefectures. Therefore, optimal policy involves increasing transfers to dense prefectures that now have even higher utility.

6.3 Comparison with real-world policy proposals

Several governments have introduced policies and policy proposals in response to concerns about the spatial consequences of depopulation. Examples include local *akiya* banks and renovation subsidies for vacant homes in Japan (Akiya Japan 2026), Italy's "1 Euro home" programs and related renovation grants (Iacubino 2025), rent and job incentives in Spain (thinkSPAIN 2021), grants for vacant or run-down island homes in Ireland (Picchi 2023), relocation payments in Switzerland (Municipality of Albinen 2017), selective island stipends and rural property-tax reductions in Greece (Giannopoulos 2026), and Scotland's Addressing Depopulation Action Plan, which funds settlement officers, rural housing, connectivity, and local recruitment strategies (Scottish Government 2024). While these policies differ in their details, they share a common goal: retain or attract residents in declining rural areas.

Relative to these proposals, the social planner that we consider chooses transfers that go in the opposite direction. Three caveats are important when comparing our results with these policies. First, we solve the planner's problem in steady state, which allows the planner to choose transfers internalizing their long-run effects on migration decisions. However, along a transition path, a planner may want to move more slowly in order to protect households who are temporarily locked into declining regions. Second, several of these policies are targeted at abandoned homes, which our model does not generate and may have negative externalities. Finally, households may value the preservation of existing communities (e.g., historic sites) in ways they do not internalize, which creates a direct rationale for transfers targeted at declining regions.

7 Conclusion

This paper studies how depopulation affects the spatial allocation of economic activity. We first document a robust pattern of scale dependence: both across and within countries, national population growth is negatively associated with changes in spatial concentration. We then show that a wide class of spatial equilibrium models generates this pattern only when housing supply elasticities fall, or housing expenditure shares rise, with density—two conditions that benchmark models assume away but that the data support. The common thread is complementarity. When land

and structures are complements in supply, and goods and housing are complements in demand, a fixed supply of land makes housing scarce and especially valuable in dense regions. Therefore, a population decline disproportionately benefits ex-ante denser regions in which housing is scarce, incentivizing migration and driving up spatial concentration in equilibrium.

Quantifying this mechanism for Japan's projected decline, we find that the long-run effects are large: Tokyo's population share quadruples from around 10% to 40%, while sparse regions such as Hokkaido nearly empty out. This reallocation of workers toward more productive regions is large enough that aggregate GDP per capita rises by 27%, even as local incomes fall everywhere. The same forces, however, make depopulation more painful for sparse regions. Therefore, a social planner would respond in the long run by increasing transfers from sparse to dense regions to encourage migration out of declining regions.

Our analysis takes the aggregate population decline as given, but spatial allocations and the decline itself may be jointly determined, which raises two questions for future work. First, the congestion and agglomeration forces that govern where households live also shape the costs of education and childcare; thus, the concentration of population into expensive cities may feed back into the aggregate fertility decline (Arkolakis et al. 2026; Borck et al. 2026). Second, housing can affect fertility directly (Lovenheim and Mumford 2013; Dettling and Kearney 2014; Fazio et al. 2025; Couillard 2026). We show that whether fertility bounces back to replacement crucially depends on the properties of housing supply. Our result is qualitative; embedding these channels into a quantitative model in which fertility, housing, and the spatial allocation are mutually determined would clarify how these forces interact and shape future depopulation.

References

Aaron, Henry J. (1966), "The social insurance paradox." *Canadian Journal of Economics and Political Science*, 32, 371–374.

Acemoglu, Daron and Pascual Restrepo (2017), "Secular stagnation? the effect of aging on economic growth in the age of automation." *American Economic Review*, 107, 174–179.

Ahlfeldt, Gabriel M. and Elisabetta Pietrostefani (2019), "The economic effects of density: A synthesis." *Journal of Urban Economics*, 111, 93–107.

Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf (2015), "The economics of density: Evidence from the berlin wall." *Econometrica*, 83, 2127–2189.

Akiya Japan (2026), "Every Akiya bank in Japan: The complete English guide." Akiya Japan.

Albouy, David, Gabriel Ehrlich, and Yingyi Liu (2016), "Housing demand, cost-of-living inequality, and the affordability crisis." NBER Working Paper 22816, National Bureau of Economic Research.

Allen, Treb and Costas Arkolakis (2014), "Trade and the topography of the spatial economy." *Quarterly Journal of Economics*, 129, 1085–1139.

Allen, Treb and Costas Arkolakis (2025), "Quantitative regional economics." In *Handbook of Regional and Urban Economics*, volume 6, 1–72, Elsevier.

Allen, Treb, Costas Arkolakis, and Xiangliang Li (2024), "On the equilibrium properties of spatial models." *American Economic Review: Insights*, 6, 472–489.

Arkolakis, Costas, Jun Hee Kwak, Jaemin Woo, and Hyunjoo Yang (2026), "Revisiting the fertility puzzle: A dynamic spatial equilibrium approach." Working paper.

Auclert, Adrien, Hannes Malmberg, Frederic Martenet, and Matthew Rognlie (2021), "Demographics, wealth, and global imbalances in the twenty-first century." Working Paper 29161, National Bureau of Economic Research.

Auerbach, Alan J., Laurence J. Kotlikoff, Robert P. Hagemann, and Giuseppe Nicoletti (1989), "The dynamics of an aging population: The case of four OECD countries." Working Paper 2797, National Bureau of Economic Research. Published as OECD Economic Studies, no. 12, pp. 97–130.

Badilla Maroto, Marco A., Benjamin Faber, Antoine B. Levy, and Mathilde Muñoz (2026), "Senior migration, local economic development and spatial concentration." NBER Working Paper 34725, National Bureau of Economic Research.

Baker, Dean, J. Bradford De Long, and Paul R. Krugman (2005), "Asset returns and economic growth." *Brookings Papers on Economic Activity*, 2005, 289–330.

Baum-Snow, Nathaniel and Gilles Duranton (2025), "Housing supply and housing affordability." In *Handbook of Regional and Urban Economics* (Dave Donaldson and Stephen J. Redding, eds.), volume 6, 353–453, Elsevier.

Baum-Snow, Nathaniel and Lu Han (2024), "The microgeography of housing supply." *Journal of Political Economy*, 132, 1897–1946.

Bloom, David E., David Canning, and Günther Fink (2010), "Implications of population ageing for economic growth." *Oxford Review of Economic Policy*, 26, 583–612.

Boppart, Timo (2014), "Structural change and the kaldor facts in a growth model with relative price effects and non-gorman preferences." *Econometrica*, 82, 2167–2196.

Borck, Rainald, Tadashi Morita, and Yasuhiro Sato (2026), "Optimal city size with endogenous fertility." Working Paper 12595, CESifo.

Caliendo, Lorenzo, Maximiliano Dvorkin, and Fernando Parro (2019), "Trade and labor market dynamics: General equilibrium analysis of the china trade shock." *Econometrica*, 87, 741–835.

Coeurdacier, Nicolas, Florian Oswald, and Marc Teignier (2025), "Structural change, land use and urban expansion." *Review of Economic Studies*.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon (2019), "The costs of agglomeration: House and land prices in french cities." *Review of Economic Studies*, 86, 1556–1589.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon (2021), "The production function for housing: Evidence from france." *Journal of Political Economy*, 129, 2766–2816.

Couillard, Benjamin K. (2026), "Build, baby, build: How housing shapes fertility." Job Market Paper, University of Toronto.

Couture, Victor, Cecile Gaubert, Jessie Handbury, and Erik Hurst (2024), "Income growth and the distributional effects of urban spatial sorting." *Review of Economic Studies*, 91, 858–898.

Cravino, Javier, Andrei A. Levchenko, and Marco Rojas (2022), "Population aging and structural transformation." *American Economic Journal: Macroeconomics*, 14, 479–498.

Cutler, David M., James M. Poterba, Louise M. Sheiner, and Lawrence H. Summers (1990), "An aging society: Opportunity or challenge?" *Brookings Papers on Economic Activity*, 1990, 1–74.

Davis, Donald R. and Jonathan I. Dingel (2019), "A spatial knowledge economy." *American Economic Review*, 109, 153–170.

Davis, Morris A. and Jonathan Heathcote (2007), "The price and quantity of residential land in the United States." *Journal of Monetary Economics*, 54, 2595–2620.

Davis, Morris A., William D. Larson, Stephen D. Oliner, and Jessica Shui (2021), "The price of residential land for counties, ZIP codes, and census tracts in the United States." *Journal of Monetary Economics*, 118, 413–431.

Desmet, Klaus and Fernando Parro (2025), "Spatial dynamics." In *Handbook of Regional and Urban Economics*, volume 6, 225–286, Elsevier.

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

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

Diamond, Rebecca and Juan Carlos Suárez Serrato (2025), "Spatial sorting and inequality." In *Handbook of Regional and Urban Economics*, volume 6, 463–513, Elsevier.

Doepke, Matthias, Anne Hannusch, Fabian Kindermann, and Michèle Tertilt (2023), "The economics of fertility: A new era." In *Handbook of the Economics of the Family*, volume 1, 151–254, Elsevier.

Eckert, Fabian, Sharat Ganapati, and Conor Walsh (2025), "Urban-biased growth: A macroeconomic analysis." Working Paper 30515, National Bureau of Economic Research.

Ehrlich, Paul R. and John P. Holdren (1971), "Impact of population growth." *Science*, 171, 1212–1217.

Fajgelbaum, Pablo D. and Cecile Gaubert (2020), "Optimal spatial policies, geography, and sorting." *Quarterly Journal of Economics*, 135, 959–1036.

Fajgelbaum, Pablo D. and Cecile Gaubert (2025), "Optimal spatial policies." In *Handbook of Regional and Urban Economics*, volume 6, 143–223, Elsevier.

Fazio, Dimas, Tarun Ramadorai, Janis Skrastins, and Bernardus Ferdinandus Van Doornik (2025), "Housing and fertility." Working paper.

Finlay, John and Trevor C. Williams (2025), "Housing demand, inequality, and spatial sorting." *Journal of International Economics*, 158, 104171.

Gaubert, Cecile and Frédéric Robert-Nicoud (2025), "Sorting to expensive cities." NBER Working Paper 33652, National Bureau of Economic Research.

Geruso, Michael and Dean Spears (2026), "The likelihood of persistently low global fertility." *Journal of Economic Perspectives*, 40, 3–26.

Giannone, Elisa, Yuhei Miyauchi, Nuno Paixão, Xinle Pang, and Yuta Suzuki (2026), "Living in a Ghost Town: The Geography of Depopulation and Aging." Working paper.

Giannopoulos, Bill (2026), "Greek island offers free food, housing, and EUR 500 monthly to attract new families." Greek City Times.

Gietel-Basten, Stuart (2016), "Japan is not the only country worrying about population decline—get used to a two-speed world." The Conversation.

Glaeser, Edward L. and Joseph Gyourko (2005), "Urban decline and durable housing." *Journal of Political Economy*, 113, 345–375.

Glaeser, Edward L., Martina Kirchberger, and Andrii Parkhomenko (2025), "Rebuilding ukraine's cities: Maximizing benefits and minimizing costs." NBER Working Paper 34598, National Bureau of Economic Research.

Gobbi, Paula E., Anne Hannusch, and Pauline Rossi (2026), "Family institutions and the global fertility transition." *Journal of Economic Perspectives*, 40, 47–70.

Greaney, Brian, Andrii Parkhomenko, and Stijn Van Nieuwerburgh (2025), "Dynamic urban economics." Working paper.

Greenwood, Michael J. (1997), "Internal migration in developed countries." In *Handbook of Population and Family Economics* (Mark R. Rosenzweig and Oded Stark, eds.), volume 1B, 647–720, North Holland.

Helpman, Elhanan (1998), "The size of regions." In *Topics in Public Economics: Theoretical and Applied Analysis* (David Pines, Efraim Sadka, and Itzhak Zilcha, eds.), 33–54, Cambridge University Press, Cambridge.

Herrendorf, Berthold, Richard Rogerson, and Ákos Valentinyi (2014), "Growth and structural transformation." In *Handbook of Economic Growth*, volume 2B, 855–941, Elsevier.

Iacubino, Carlo (2025), "Sardinia depopulation incentives 2025: EUR 600 newborn bonus, housing grants, and business support." Idealista News.

Jones, Charles I. (2022), "The end of economic growth? unintended consequences of a declining population." *American Economic Review*, 112, 3489–3527.

Kleinman, Benny, Ernest Liu, and Stephen J. Redding (2023), "Dynamic spatial general equilibrium." *Econometrica*, 91, 385–424.

Krugman, Paul (1991), "Increasing returns and economic geography." *Journal of Political Economy*, 99, 483–499.

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

Lucas, Robert E. B. (1997), "Internal migration in developing countries." In *Handbook of Population and Family Economics* (Mark R. Rosenzweig and Oded Stark, eds.), volume 1B, 721–798, North Holland.

Mankiw, N. Gregory and David N. Weil (1989), "The baby boom, the baby bust, and the housing market." *Regional Science and Urban Economics*, 19, 235–258.

Michaels, Guy, Ferdinand Rauch, and Stephen J. Redding (2012), "Urbanization and structural transformation." *Quarterly Journal of Economics*, 127, 535–586.

Mori, Tomoya and Daisuke Murakami (2025), "Sustainability of Cities under Declining Population and Decreasing Distance Frictions: The Case of Japan." *arXiv preprint arXiv:2505.08333*.

Morita, Hiroshi (2020), "Fiscal multipliers in the most aged country: Empirical evidence and theoretical interpretation." HIAS Discussion Paper No. HIAS-E-100, Hitotsubashi Institute for Advanced Study, Hitotsubashi University.

Muellbauer, John (1975), "Aggregation, income distribution and consumer demand." *Review of Economic Studies*, 42, 525–543.

Muellbauer, John (1976), "Community preferences and the representative consumer." *Econometrica*, 44, 979–999.

Municipality of Albinen (2017), "Declaration of the municipality of Albinen (Switzerland/Wallis) on housing subsidies and support for families." Municipality of Albinen.

National Institute of Population and Social Security Research (2006), "Population projections for Japan (December 2006): 2006–2055." Technical report, National Institute of Population and Social Security Research, Tokyo.

Pesaresi, Martino, Marcello Schiavina, Panagiotis Politis, Sergio Freire, Katarzyna Krasnodebska, Johannes H. Uhl, Alessandra Carioli, Christina Corbane, Lewis Dijkstra, Pietro Florio, Hannah K. Friedrich, Jing Gao, Stefan Leyk, Linlin Lu, Luca Maffenini, Ines Mari-Rivero, Michele Melchiorri, Vasileios Syrris, Jamon Van Den Hoek, and Thomas Kemper (2024), "Advances on the global human settlement layer by joint assessment of earth observation and population survey data." *International Journal of Digital Earth*, 17, 2390454.

Peters, Michael and Conor Walsh (2026), "Population growth and firm dynamics." *Journal of Political Economy: Macroeconomics*, 4.

Piazzesi, Monika and Martin Schneider (2016), "Housing and macroeconomics." In *Handbook of Macroeconomics*, volume 2, 1547–1640, Elsevier.

Picchi, Aimee (2023), "Ireland is paying up to $92,000 to people who buy homes on remote islands. here's how it works." CBS News.

Poterba, James M. (2001), "Demographic structure and asset returns." *Review of Economics and Statistics*, 83, 565–584.

Redding, Stephen J. (2016), "Goods trade, factor mobility and welfare." *Journal of International Economics*, 101, 148–167.

Redding, Stephen J. (2025), "Quantitative urban economics." In *Handbook of Regional and Urban Economics*, volume 6, 73–141, Elsevier.

Redding, Stephen J. and Esteban Rossi-Hansberg (2017), "Quantitative spatial economics." *Annual Review of Economics*, 9, 21–58.

Roback, Jennifer (1982), "Wages, rents, and the quality of life." *Journal of Political Economy*, 90, 1257–1278.

Robert-Nicoud, Frédéric, Pierre-Philippe Combes, Gilles Duranton, and Laurent Gobillon (2026), "A unified urban model with non-homothetic housing demand." NBER Working Paper 34876, National Bureau of Economic Research.

Rognlie, Matthew (2015), "Deciphering the fall and rise in the net capital share: Accumulation or scarcity?" *Brookings Papers on Economic Activity*, 2015, 1–69.

Rosen, Sherwin (1979), "Wage-based indexes of urban quality of life." In *Current Issues in Urban Economics* (Peter Mieszkowski and Mahlon Straszheim, eds.), 74–104, Johns Hopkins University Press, Baltimore.

Saiz, Albert (2010), "The geographic determinants of housing supply." *Quarterly Journal of Economics*, 125, 1253–1296.

Samuelson, Paul A. (1958), "An exact consumption-loan model of interest with or without the social contrivance of money." *Journal of Political Economy*, 66, 467–482.

Schiavina, Marcello, Sergio Freire, Alessandra Carioli, and Kytt MacManus (2023), "GHS-POP R2023A - GHS population grid multitemporal (1975-2030)." European Commission, Joint Research Centre dataset.

Scottish Government (2024), "Supporting and enabling sustainable communities: Action plan to address depopulation." Scottish Government.

thinkSPAIN (2021), "Teruel village offers free rent and job for family with young children." thinkSPAIN.

Weil, David N. (2026), "How much would continued low fertility affect the US standard of living?" *Journal of Economic Perspectives*, 40, 27–46.

Online appendix for

Making room on an empty planet: The spatial consequences of depopulation

Tim de Silva
Stanford GSB

James D. Paron
Stanford GSB

June 17, 2026

Latest Version

Contents

Appendix A. Additional empirical results

Figure A.1. Map of Japanese prefectures

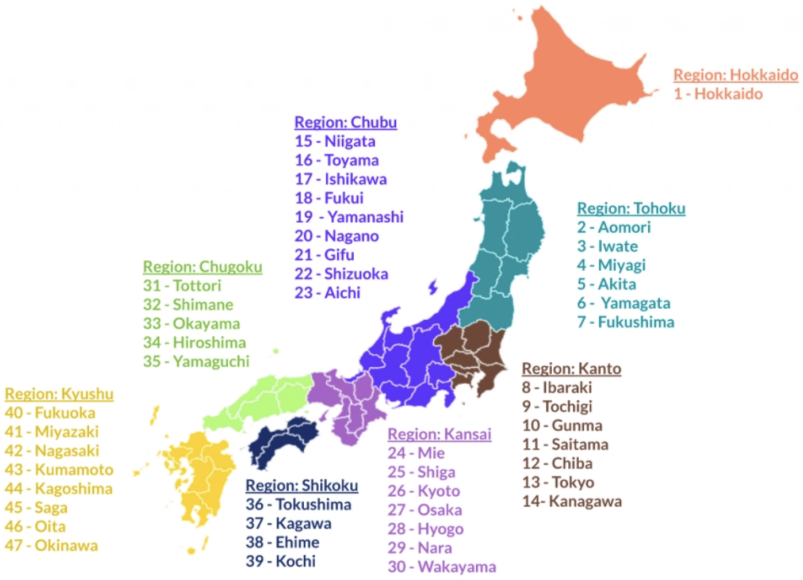

Notes: This figure shows a map of Japanese prefectures taken from here. Go back.

Figure A.2. Tokyo population share and net migration

Panel A: Tokyo population share and migration rate

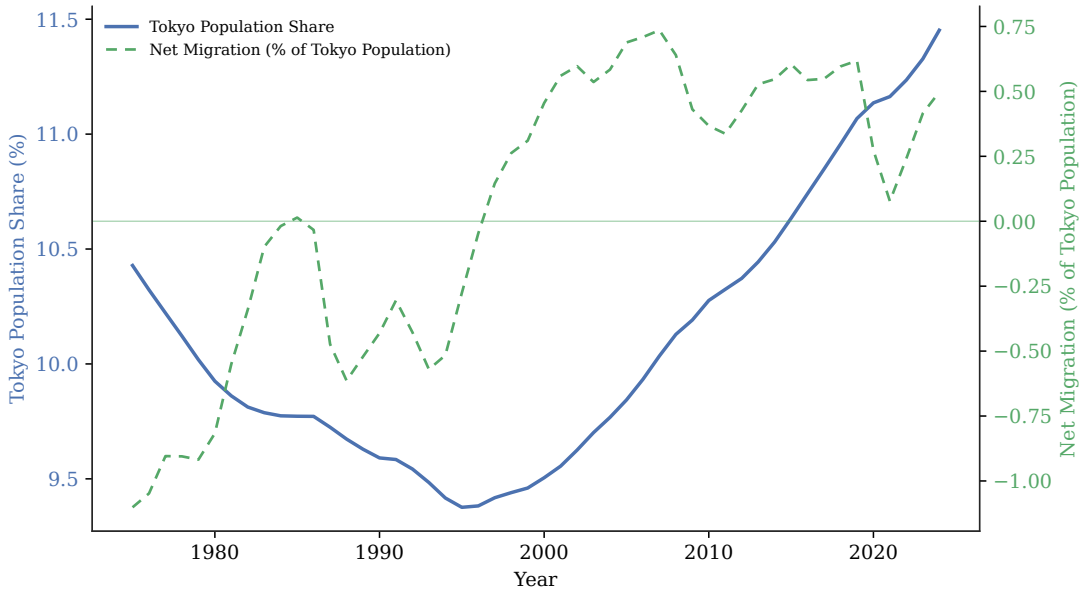

Panel B: Contribution of migration to Tokyo population share

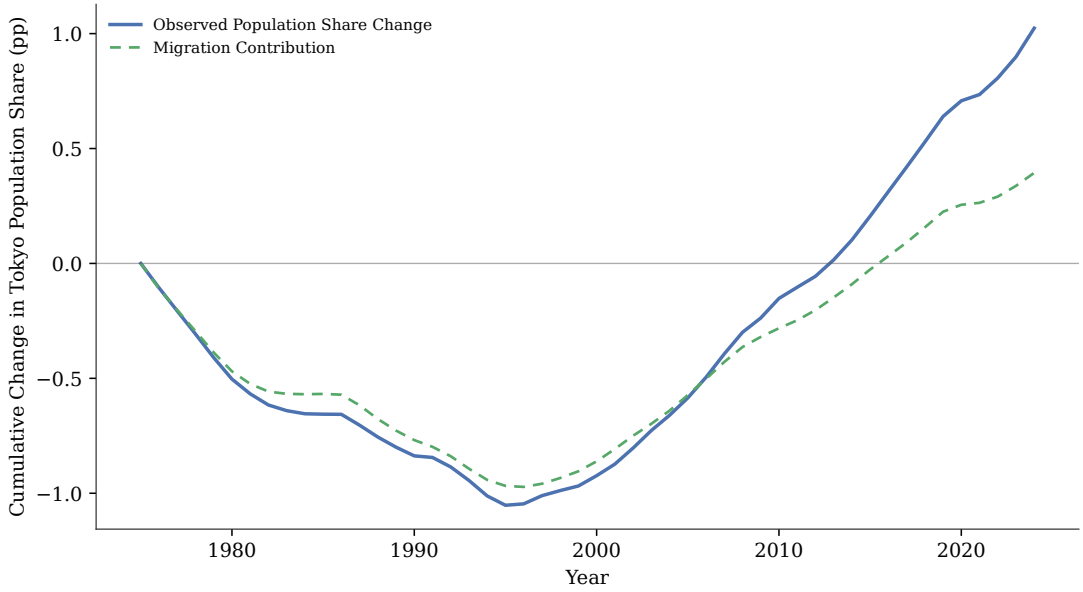

Notes: The figure uses annual Tokyo population, in-migrants, and out-migrants from the Statistics Bureau of Japan e-Stat Statistics Dashboard. Panel A shows Tokyo's population share and net migration into Tokyo, defined as in-migrants minus out-migrants divided by Tokyo population multiplied by 100. Panel B shows the contribution of migration to Tokyo's population share starting in 1975, computed by cumulating annual Tokyo net migration divided by lagged total Japan population. Go back.

Figure A.3. Internal in-migration by age: Tokyo and other prefectures

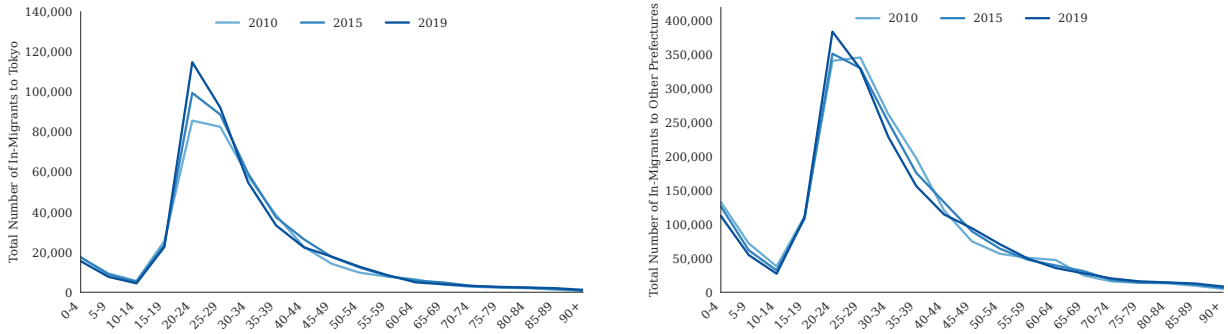

Notes: The figure uses the e-Stat Report on Internal Migration in Japan age-profile table. The left panel plots Japanese, both-sex in-migrants from other prefectures and other cities to Tokyo by five-year age group in 2010, 2015, and 2019. The right panel plots the same measure for all non-Tokyo destinations, computed as the Japan aggregate minus Tokyo. The final age group contains ages 90 and older. Go back.

Figure A.4. Cumulative net migration to Tokyo by origin prefecture, 2010–2019

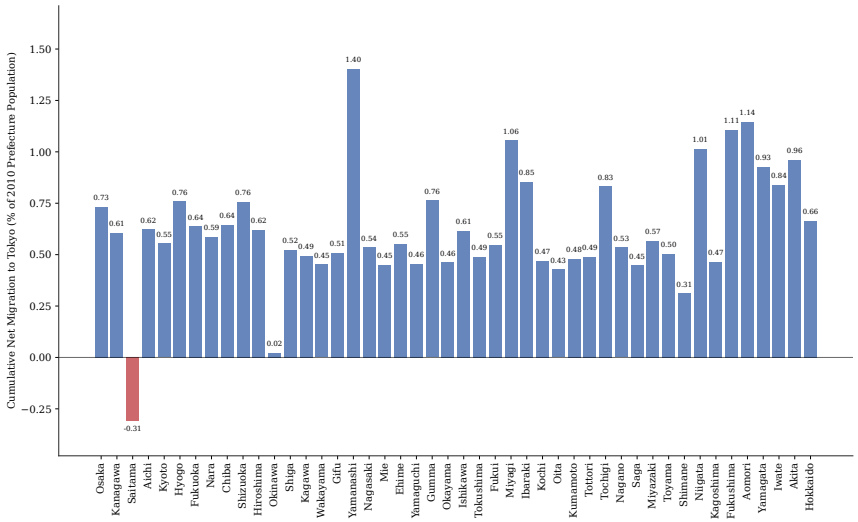

Notes: The figure uses the e-Stat Report on Internal Migration in Japan prefecture origin-destination table. For each non-Tokyo prefecture, the bar sums Japanese, both-sex migrants from that prefecture to Tokyo minus migrants from Tokyo to that prefecture over 2010–2019, divides by the prefecture's 2010 population from the e-Stat Statistics Dashboard, and multiplies by 100. Prefectures are sorted by 2000 population share over fixed-2010 inhabitable land share, matching Figure 2. Go back.

Figure A.5. Comparing GHSL and official Japanese population data

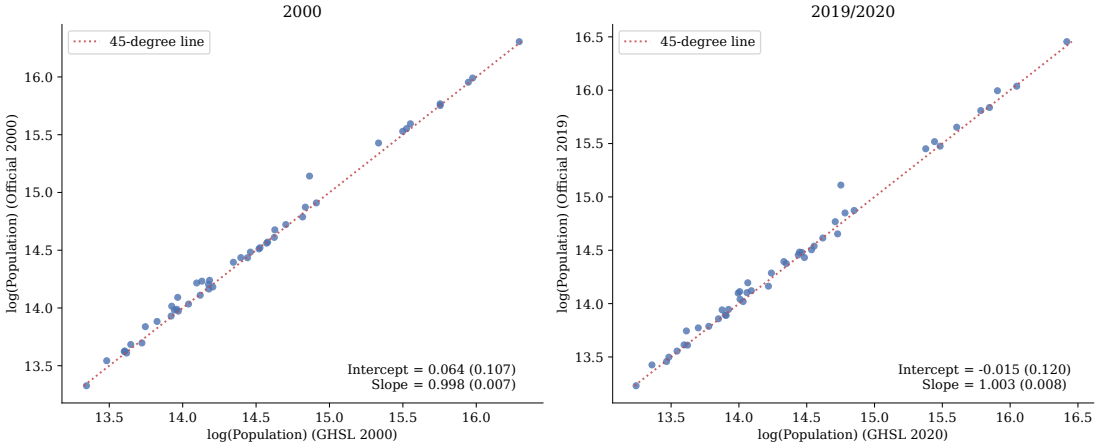

Notes: This figure compares prefecture-level population counts in Japan between the data from the Statistics Bureau of Japan used in Section 1.1 and the GHSL data used in Section 1.2. The x-axis in both plots is log GHSL population; the y-axis in both plots is log Statistics Bureau of Japan population. Both panels report intercepts and slope coefficients from unweighted regressions with the corresponding heteroskedasticity-robust standard errors. The left panel compares data in 2000; the right panel compares data in 2019 and 2020 such that both sets of points represent the most recent values before the start of COVID-19. Go back.

Figure A.6. Population growth and spatial concentration across countries, by continent

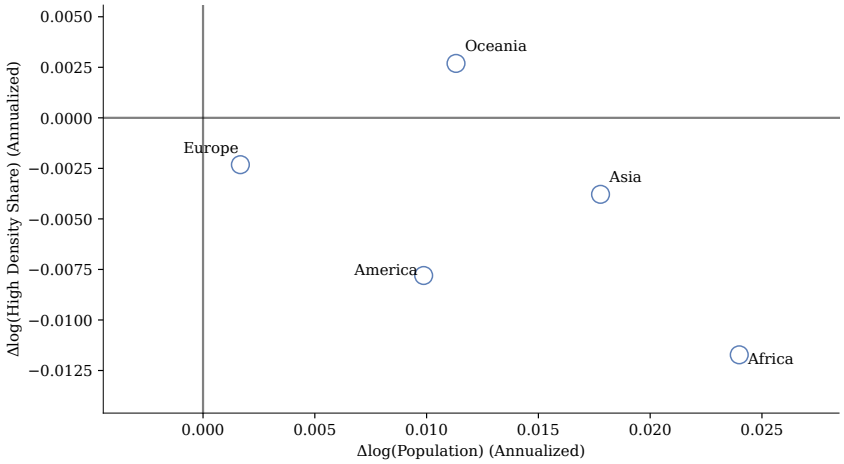

Notes: This figure reproduces Figure 3, averaging the annualized population-growth and high-density-share variables by continent. Go back.

Table A.1. Cross-country evidence of scale dependence: continent fixed effects

	$\Delta \log$(High Density Share)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta \log$(Population)	-0.161	-0.156	-0.175	-0.169	-0.476	-0.320	-0.384	-0.300
	(0.071)	(0.075)	(0.079)	(0.086)	(0.121)	(0.109)	(0.176)	(0.157)
$\Delta \log$(Agriculture Share)		-0.092		-0.093		-0.155		-0.151
		(0.078)		(0.078)		(0.071)		(0.079)
$\Delta \log$(Old-Age Dependency Ratio)			-0.035	-0.032			0.081	0.020
			(0.063)	(0.072)			(0.102)	(0.110)
Constant	-0.008	-0.009	-0.008	-0.009	-0.002	-0.008	-0.004	-0.008
	(0.002)	(0.003)	(0.002)	(0.003)	(0.004)	(0.004)	(0.004)	(0.004)
2000 Population Weights	No	No	No	No	Yes	Yes	Yes	Yes
Continent Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	190	178	190	178	190	178	190	178
R^2	0.205	0.211	0.206	0.212	0.307	0.342	0.311	0.342

Notes: This table reproduces Table 1 with continent fixed effects. Go back.

Table A.2. Cross-country evidence of scale dependence: controlling for GDP growth

	$\Delta \log$(High Density Share)			
	(1)	(2)	(3)	(4)
$\Delta \log$(Population)	-0.235	-0.181	-0.544	-0.328
	(0.068)	(0.077)	(0.123)	(0.102)
$\Delta \log$(Agriculture Share)		-0.095		-0.214
		(0.073)		(0.079)
$\Delta \log$(Old-Age Dependency Ratio)		0.038		0.057
		(0.074)		(0.067)
$\Delta \log$(GDP Per Capita)	-0.018	-0.017	-0.029	-0.059
	(0.023)	(0.024)	(0.020)	(0.022)
Constant	-0.002	-0.005	0.002	-0.004
	(0.002)	(0.002)	(0.002)	(0.002)
2000 Population Weights	No	No	Yes	Yes
Observations	186	174	186	174
R^2	0.092	0.118	0.325	0.390

Notes: This table reproduces the unweighted and 2000-population-weighted specifications from Table 1 with annualized log GDP per capita growth over the same time period as a control. Go back.

Table A.3. Cross-country evidence of scale dependence: raw change in high-density share

	Δ High Density Share							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta \log$(Population)	-0.032	-0.026	-0.027	-0.022	-0.078	-0.061	-0.070	-0.058
	(0.010)	(0.011)	(0.012)	(0.013)	(0.014)	(0.015)	(0.016)	(0.016)
$\Delta \log$(Agriculture Share)		-0.015		-0.015		-0.018		-0.018
		(0.009)		(0.009)		(0.010)		(0.011)
$\Delta \log$(Old-Age Dependency Ratio)			0.009	0.008			0.009	0.004
			(0.012)	(0.013)			(0.010)	(0.012)
Constant	-0.001	-0.001	-0.001	-0.001	-0.000	-0.001	-0.000	-0.001
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
2000 Population Weights	No	No	No	No	Yes	Yes	Yes	Yes
Observations	190	178	190	178	190	178	190	178
R^2	0.062	0.080	0.065	0.083	0.320	0.346	0.324	0.347

Notes: This table reproduces Table 1 using annualized raw changes in the high-density share rather than annualized log changes. Go back.

Figure A.7. Population growth and spatial concentration across countries: land-density measure

Panel A: Raw change

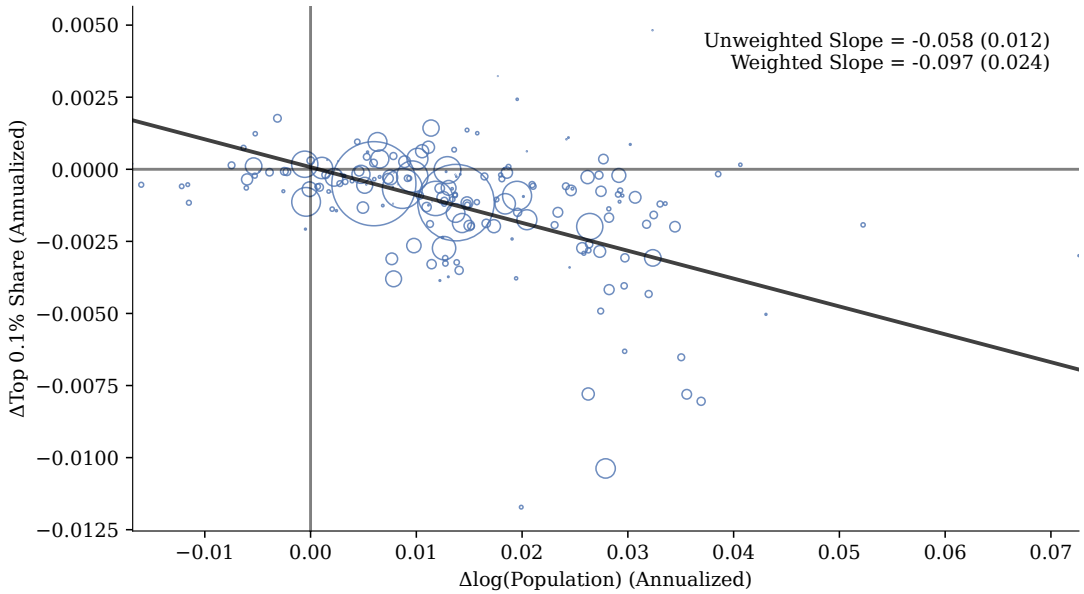

Panel B: Log change

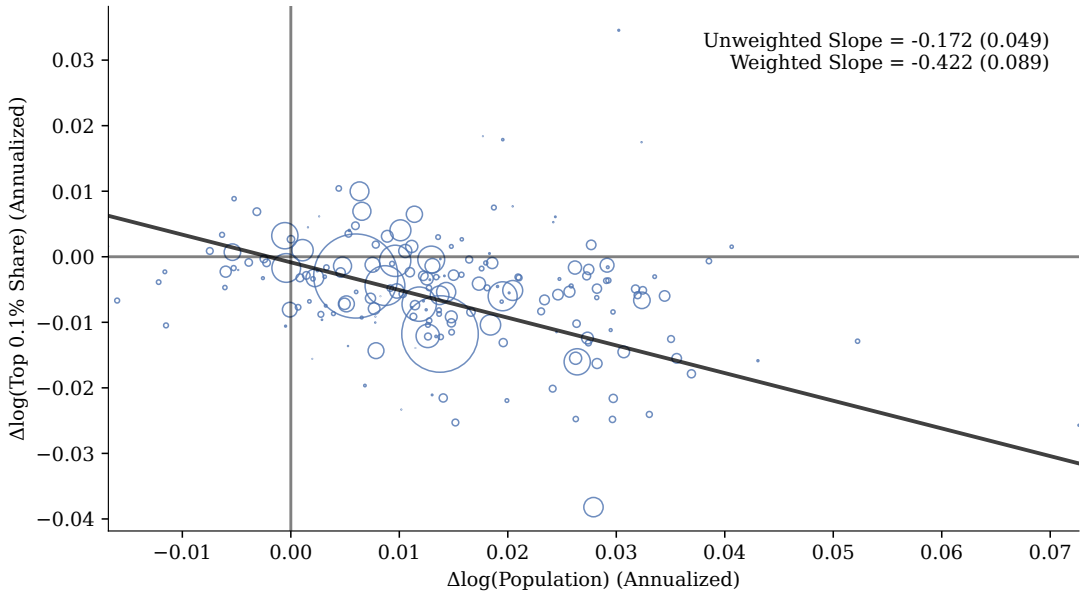

Notes: This figure reproduces Figure 3 using an alternative concentration measure: the share of national population living in the country's 2000 top 0.1% of grid cells by population count. Panel A uses annualized raw changes in this share, while Panel B uses annualized log changes. Go back.

Figure A.8. Land share of property values and population density

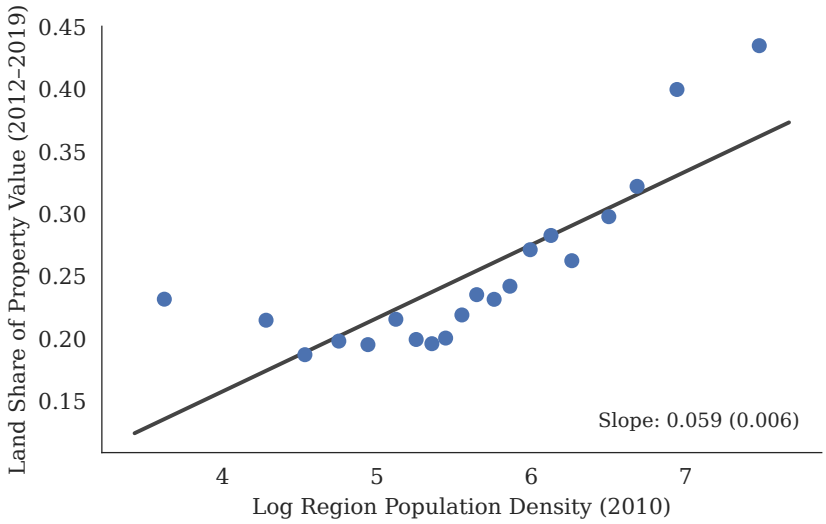

Notes: This figure uses the CTPP-region construction described in the notes to Figure 4. The vertical axis is the 2010-population-weighted average across counties in each CTPP region of the 2012–2019 county-level land share of property value from Davis et al. (2021). The horizontal axis is log 2010 population density, computed as 2010 Census SF1 county population divided by county-gazetteer land area, with both summed over counties assigned to each region through the BSH tract-to-region mapping. Dots are averages within equally-spaced bins; the fitted line and reported slope are estimated on the unbinned region-level data using OLS with heteroskedasticity-robust standard errors in parentheses. Go back.

Appendix B. Proofs of results in main text

B.1 Proof of Proposition 1

The first statement follows immediately from (8). The population changes $d\log N_i/d\log N$ are identical across i (scale independence) if and only if $\varepsilon_i^h(N_i)$ are identical across i.

Now consider the second statement. Identical elasticities for all primitives means that there exists a constant ν such that $\varepsilon_i^h(N_i) = \nu$ for all parameters $\{\{A_i, L_i\}_i, N, \alpha, \sigma, \theta\}$ that admit a unique and interior equilibrium (interior meaning $N_i > 0\ \forall i$). That is,

$$\nu = -\frac{\partial \log h_i^{\mathsf{S}}(N_i)}{\partial \log N_i}, \tag{B.1}$$

for some ν. The solution to this differential equation is the per-capita housing supply function

$$h_i^{\mathsf{S}}(N_i) = \bar{h}_i N_i^{-\nu} \tag{B.2}$$

for some constant $\bar{h}_i$. To express this as total housing supply, use $H_i^{\mathsf{S}}(X_i, L_i) = N_i h_i^{\mathsf{S}}(N_i)$ and $N_i = X_i/A_X$ to get

$$H_i^{\mathsf{S}}(X_i, L_i) = \bar{h}_i N_i^{1-\nu} = \frac{\bar{h}_i}{A_X^{1-\nu}} X_i^{1-\nu}. \tag{B.3}$$

Defining functions $G_i(L_i) \equiv \bar{h}_i / A_X^{1-\nu}$ (which are constants, since L_i is fixed) gives the result.

B.2 Proof of Proposition 2

The first statement is a corollary of Proposition 1. To prove the second statement, we first show a useful lemma. Recall that "spatial concentration" is defined as the share of population living in the $\bar{J} < J$ densest places. Re-indexing locations $j \in \{1, \ldots, J\}$ in order of density,

$$\frac{N_1}{L_1} \geq \frac{N_2}{L_2} \geq \cdots \geq \frac{N_J}{L_J} \qquad \text{with} \qquad \frac{N_{\bar{J}}}{L_{\bar{J}}} > \frac{N_{\bar{J}+1}}{L_{\bar{J}+1}} \quad \text{(no ties)}, \tag{B.4}$$

spatial concentration is therefore defined as

$$S_{\bar{J}} \equiv \sum_{j=1}^{\bar{J}} \frac{N_j}{N}. \tag{B.5}$$

Thus, the change in spatial concentration in response to a population change is

$$\frac{d\log S_{\bar{J}}}{d\log N} = \sum_{j=1}^{\bar{J}} \frac{N_j}{\sum_{j'=1}^{\bar{J}} N_{j'}} \left(\frac{d\log N_j}{d\log N} - 1 \right). \tag{B.6}$$

From this we get the lemma.

Lemma 1. *If the population response $d\log N_j/d\log N$ is strictly falling in density N_j/L_j, then $d\log S_{\bar{J}}/d\log N < 0$: spatial concentration falls (rises) when N rises (falls). If the response $d\log N_j/d\log N$ is strictly rising in density, then $d\log S_{\bar{J}}/d\log N > 0$.*

Proof. Because population levels N_j must sum to N, it follows that

$$1 = \sum_{j=1}^{J} \frac{dN_j}{dN} = \sum_{j=1}^{J} \frac{N_j}{N} \frac{d\log N_j}{d\log N}. \tag{B.7}$$

The weighted average of $d\log N_j/d\log N$ is one. Hence, if there are any differences in $d\log N_j/d\log N$, then $\min_j d\log N_j/d\log N < 1$ and $\max_j d\log N_j/d\log N > 1$. No ties ensures there are differences. Consequently, any weighted average of $d\log N_j/d\log N$ that excludes the largest values will necessarily be less than one. In particular, if $d\log N_j/d\log N$ is strictly falling in density N_j/L_j, then the change in spatial concentration (B.6) contains such a weighted average, and so it must be that $d\log S_{\bar{J}}/d\log N < 0$. If instead $d\log N_j/d\log N$ is strictly rising in density N_j/L_j, then the same reasoning implies $d\log S_{\bar{J}}/d\log N > 0$. □

Now, from (8), $d\log N_j/d\log N$ is falling in the elasticities $\varepsilon_i^h(N_i)$. For $\eta < 1$, these elasticities are increasing in density N_i/L_i (and only in density), so it follows from the lemma that $d\log S_{\bar{J}}/d\log N < 0$. If instead $\eta > 1$, then the lemma implies $d\log S_{\bar{J}}/d\log N > 0$.

B.3 Proof of Proposition 3

To prove the first statement, we need to first determine $d\log p_i/d\log N_i$. To do this, we will totally differentiate the market-clearing conditions $h_i^{\mathrm{S}}(w_i, p_i, N_i) = h_i^{\mathrm{D}}(w_i, p_i)$. The demand change is

$$\frac{d\log h_i^{\mathrm{D}}}{d\log N_i} = \frac{\partial \log h_i^{\mathrm{D}}}{\partial \log p_i} \frac{d\log p_i}{d\log N_i} + \frac{\partial \log h_i^{\mathrm{D}}}{\partial \log w_i} \frac{d\log w_i}{d\log N_i} \tag{B.8}$$

$$= -\zeta_i^{\mathrm{D}} \frac{d\log p_i}{d\log N_i} + \sigma. \tag{B.9}$$

The second equality substitutes in the definition of the price elasticity of demand and uses (13) (i.e., an income elasticity of one) and $w_i = A_i N_i^\sigma$. The supply change is

$$\frac{d\log h_i^{\mathrm{S}}}{d\log N_i} = \frac{H_i'(X_i)X_i}{H_i(X_i)}\frac{d\log X_i}{d\log N_i} - 1. \tag{B.10}$$

To express this in terms of price changes, differentiate the construction sector's first-order condition (15) to get

$$\frac{d\log w_i}{d\log N_i} = \frac{d\log p_i}{d\log N_i} + \frac{H_i''(X_i)X_i}{H_i'(X_i)}\frac{d\log X_i}{d\log N_i}. \tag{B.11}$$

and substitute this back in to get[1]

$$\frac{d\log h_i^{\mathrm{S}}}{d\log N_i} = -\frac{H_i'(X_i)}{H_i(X_i)}\frac{H_i'(X_i)}{H_i''(X_i)}\left(\frac{d\log p_i}{d\log N_i} - \frac{d\log w_i}{d\log N_i}\right) - 1 = \zeta_i^{\mathrm{S}}\left(\frac{d\log p_i}{d\log N_i} - \sigma\right) - 1. \tag{B.12}$$

Setting change in supply equal to change in demand and solving for the price change, we have

$$\frac{d\log p_i}{d\log N_i} = \frac{1}{\zeta_i^{\mathrm{S}} + \zeta_i^{\mathrm{D}}} + \frac{\zeta_i^{\mathrm{S}} + 1}{\zeta_i^{\mathrm{S}} + \zeta_i^{\mathrm{D}}}\sigma. \tag{B.13}$$

Substituting this into (16), we get (17).

For the second statement, we require that the objects $p_i h_i / w_i$, ζ_i^{D}, and ζ_i^{S} are constant and equal across locations for all parameters $\{\{A_i, L_i\}_i, N, \sigma, \theta\}$ that admit a unique and interior equilibrium (interior meaning $N_i > 0\ \forall i$). Start with ζ_i^{S}. Identical supply elasticities means that there exists a constant $\tilde{\nu}$ such that $\zeta_i^{\mathrm{S}} = \tilde{\nu}$ for all i:

$$\tilde{\nu} = -\frac{H_i'(X_i)}{H_i(X_i)}\frac{H_i'(X_i)}{H_i''(X_i)}, \tag{B.14}$$

for some $\tilde{\nu}$. The solution to this differential equation is

$$H_i(X_i) = \bar{H}_{1i}(\bar{H}_{2i} + X_i)^{\frac{\tilde{\nu}}{1+\tilde{\nu}}}, \tag{B.15}$$

for some constants $\bar{H}_{1i}$ and $\bar{H}_{2i}$. In order to satisfy the inada condition $\lim_{X_i \downarrow 0} H_i'(X_i) = \infty$, we must have $\bar{H}_{2i} = 0$. Thus, defining functions $G_i(L_i) \equiv \bar{H}_{1i}$ and constant $\nu \equiv 1/(1+\tilde{\nu})$, we get the functional form stated in the proposition.

Now consider demand elasticities ζ_i^{D}. To get $\zeta_i^{\mathrm{D}} = \zeta_p$ (a constant) for all i, indirect utility $v(p_i)$

[1] This is why ζ_i^{S} can be interpreted as the price elasticity: it is the change in per-capita housing supply in response to a change in house prices, holding wages fixed.

must solve

$$\zeta_p = -p_i \left(\frac{v''(p_i)}{v'(p_i)} - \frac{v'(p_i)}{v(p_i)} \right). \tag{B.16}$$

For $\zeta_p \neq 1$, the solution to this is

$$v(p_i) = \bar{v} \exp\left\{ -\alpha \frac{p_i^{1-\zeta_p} - 1}{1 - \zeta_p} \right\}. \tag{B.17}$$

The limiting case as $\zeta_p \to 1$ is $v(p_i) = \bar{v} p_i^{-\alpha}$. Moreover, by Roy's identity, the housing expenditure share implied by this is

$$\frac{p_i h_i}{w_i} = -\frac{\partial \log v(p_i)}{\partial \log p_i} = \alpha p_i^{1-\zeta_p}. \tag{B.18}$$

This expenditure share is constant for all p_i if and only if $\zeta_p = 1$ or equivalently $v(p_i) \propto p_i^{-\alpha}$. This indirect utility corresponds to Cobb-Douglas utility $U(c_i, h_i) = \bar{U} c_i^{1-\alpha} h_i^{\alpha}$ (up to a positive constant $\bar{U}$). This proves the proposition.

B.4 Proof of Proposition 4

The first statement is an application of Proposition 3. To show the second statement, we can apply Lemma 1 from the proof of Proposition 2. In order to apply that lemma, we must show that $d \log N_j / d \log N$ is strictly falling in density (equivalently, in population level, since we have chosen $L_i = L$). From (17) with (20) and (22), we see that this specifically means that it suffices to show that X_i and p_i are rising in N_i.

The idea of this proof is to find an equilibrium locus $(X(N_i), p(N_i))$ and show that both outcomes are higher in locations with higher equilibrium N_i. Starting with the population shares (14), there is a constant Ω, common across locations, such that

$$\left(\frac{N_i}{N} \right)^{\theta - \sigma} = \Omega A_i v(p_i). \tag{B.19}$$

Substituting in $A_i = w_i / N_i^{\sigma}$, then $w_i = p_i H'(X_i) A_X$, this implies

$$N_i^{\theta} = \Omega A_X p_i H'(X_i) v(p_i) N^{\theta - \sigma}. \tag{B.20}$$

Moreover, equating housing supply $h_i^{\mathrm{S}} = H(X_i)/N_i$ with housing demand $h_i^{\mathrm{D}} = \alpha w_i p_i^{-\zeta_p}$ (and again using $w_i = p_i H'(X_i) A_X$), we have

$$N_i = \frac{H(X_i)}{\alpha A_X H'(X_i) p_i^{1-\zeta_p}}. \tag{B.21}$$

The conditions (B.20) and (B.21) imply that all locations lie on the same equilibrium locus $(X(N_i), p(N_i))$. We now prove that this locus is increasing in N_i. Differentiating (B.20) and (B.21) across population levels N_i gives the system[2]

$$\theta = \left(1 - \frac{p_i h_i}{w_i}\right) \frac{d \log p_i}{d \log N_i} + \frac{H''(X_i) X_i}{H'(X_i)} \frac{d \log X_i}{d \log N_i} \tag{B.22}$$

$$1 = -(1 - \zeta_p) \frac{d \log p_i}{d \log N_i} + \left(\frac{H'(X_i) X_i}{H(X_i)} - \frac{H''(X_i) X_i}{H'(X_i)}\right) \frac{d \log X_i}{d \log N_i}. \tag{B.23}$$

where we have used $v'(p_i) p_i / v(p_i) = -p_i h_i / w_i$. The solution to this system is

$$\frac{d \log p_i}{d \log N_i} = \frac{1}{K_i} \left(\theta \frac{H'(X_i) X_i}{H(X_i)} - (1 + \theta) \frac{H''(X_i) X_i}{H'(X_i)}\right) > 0, \tag{B.24}$$

$$\frac{d \log X_i}{d \log N_i} = \frac{1}{K_i} \left(1 - \frac{p_i h_i}{w_i} + \theta(1 - \zeta_p)\right) > 0, \tag{B.25}$$

where the denominator is

$$K_i \equiv \left(1 - \frac{p_i h_i}{w_i}\right) \frac{H'(X_i) X_i}{H(X_i)} - \left(\zeta_p - \frac{p_i h_i}{w_i}\right) \frac{H''(X_i) X_i}{H'(X_i)} > 0. \tag{B.26}$$

These are positive since we have assumed that $H(X_i), H'(X_i) > 0$, $H''(X_i) < 0$, and $\zeta_p > p_i h_i / w_i$ (recall that this is required for positive Hicksian substitution elasticities). The fact that these derivatives are positive implies that, for any locations i and j with $N_i > N_j$ in equilibrium, it must be that $p_i > p_j$ and $X_i > X_j$.

Since $\eta \leq 1$, (20) implies that ζ_i^S is weakly decreasing in X_i / L, and hence weakly decreasing in N_i. Since $\zeta_p \leq 1$, (22) implies that the housing expenditure share $p_i h_i / w_i = \alpha p_i^{1 - \zeta_p}$ is weakly increasing in p_i, and hence weakly increasing in N_i. Therefore the denominator in (23) is weakly increasing in N_i, and strictly increasing if at least one of $\eta < 1$ or $\zeta_p < 1$ holds. Hence $d \log N_i / d \log N$ is strictly falling in density. The result then follows from Lemma 1, by the exact same reasoning as in the proof of Proposition 2.

B.5 Proof of Proposition 5

In order to simplify the proof, we solve a slightly simpler problem and then show from the solution to that problem that (47) solves the original problem. In particular, we start from the assumption that average local profits $\pi_i = \Pi_i / N_i$, not national profits $\bar{\pi}$, are redistributed to local residents. Thus, the budget constraint is

$$x_i \equiv w_i + \pi_i + \tau_i \geq c_i + p_i h_i, \tag{B.27}$$

[2]This is a derivative along the locus, and hence a comparison across different locations in equilibrium. It is not the local change in p_i and X_i in response to a local population change.

where x_i is total expenditure. The consequence of this assumption is that local house prices and wages will be functions only of local population and transfers—$w_i(N_i, \tau_i)$ and $p_i(N_i, \tau_i)$—instead of the whole distribution $\{N_j, \tau_j\}_j$.

An important feature of the planner's problem is that changing the transfer τ_j affects utility not only directly through redistribution of consumption, but also indirectly by causing migration. The first step in solving the planner's problem is thus to express the objective function in terms of both τ_j and N_j. Note first that, conditional on choosing location i, a household's utility (before taste shocks) corresponds to indirect utility $U(c_i, h_i) = U^*(x_i, p_i) = x_i v(p_i)$, which in equilibrium can be written as some function

$$f_i(N_i, \tau_i) \equiv U^*(x_i(N_i, \tau_i), p_i(N_i, \tau_i)). \tag{B.28}$$

Thus, the planner's objective can be written as

$$\mathcal{U} \equiv \max_{\{\tau_i\}} \mathbb{E}\left[\max_i \{\log(f_i(N_i, \tau_i)\epsilon_i)\}\right], \tag{B.29}$$

subject to incentive compatibility, which requires that the shares obey

$$\frac{N_i}{N} = \frac{(x_i v(p_i))^{\frac{1}{\theta}}}{\sum_j (x_j v(p_j))^{\frac{1}{\theta}}} = \frac{f_i(N_i, \tau_i)^{\frac{1}{\theta}}}{\sum_j f_j(N_j, \tau_j)^{\frac{1}{\theta}}}. \tag{B.30}$$

By the properties of the Fréchet distribution (Gumbel in logs), the objective is

$$\mathcal{U} = \max_{\{\tau_j\}} \theta \log\left(\sum_i f_i(\tau_i, N_i)^{\frac{1}{\theta}}\right) + \theta\bar{\gamma}. \tag{B.31}$$

The planner's Lagrangian is then

$$\mathcal{L} = \mathcal{U} - \lambda \sum_j \tau_j N_j, \tag{B.32}$$

implying first-order conditions for each τ_i

$$\frac{d\mathcal{U}}{d\tau_i} = \lambda\left(N_i + \sum_j \tau_j \frac{dN_j}{d\tau_i}\right). \tag{B.33}$$

Using the population shares, the left-hand side evaluates to

$$\frac{d\mathcal{U}}{d\tau_i} = \sum_j \frac{N_j}{N}\left(\frac{\partial \log f_j(\tau_j, N_j)}{\partial \tau_i} + \frac{\partial \log f_j(\tau_j, N_j)}{\partial \log N_j}\frac{d \log N_j}{d\tau_i}\right) \tag{B.34}$$

In order to solve for $\{\tau_j\}$, we must then determine the first-order equilibrium response of each N_j to

each τ_i. Differentiating the population shares (B.30) with respect to τ_i gives

$$\frac{d\log N_j}{d\tau_i} = \frac{1}{\theta}\left(\frac{\partial \log f_j}{\partial \tau_i} + \frac{\partial \log f_j}{\partial \log N_j}\frac{d\log N_j}{d\tau_i}\right) - \frac{1}{\theta}\sum_{j'}\frac{N_{j'}}{N}\left(\frac{\partial \log f_{j'}}{\partial \tau_i} + \frac{\partial \log f_{j'}}{\partial \log N_{j'}}\frac{d\log N_{j'}}{d\tau_i}\right) \tag{B.35}$$

$$= \left(\theta - \frac{\partial \log f_j}{\partial \log N_j}\right)^{-1}\left(\frac{\partial \log f_j}{\partial \tau_i} - \frac{d\mathcal{U}}{d\tau_i}\right). \tag{B.36}$$

Given this solution, it follows that

$$\frac{d\mathcal{U}}{d\tau_i} = \sum_j \frac{N_j}{N}\left(\frac{\partial \log f_j}{\partial \tau_i} + \frac{\partial \log f_j}{\partial \log N_j}\frac{d\log N_j}{d\tau_i}\right) \tag{B.37}$$

$$= \sum_j \frac{N_j}{N}\left(\frac{\partial \log f_j}{\partial \tau_i} + \frac{\partial \log f_j}{\partial \log N_j}\left(\theta - \frac{\partial \log f_j}{\partial \log N_j}\right)^{-1}\left(\frac{\partial \log f_j}{\partial \tau_i} - \frac{d\mathcal{U}}{d\tau_i}\right)\right) \tag{B.38}$$

$$= \left(\sum_j \frac{N_j}{N}\left(\theta - \frac{\partial \log f_j}{\partial \log N_j}\right)^{-1}\right)^{-1}\left(\sum_j \frac{N_j}{N}\left(\theta - \frac{\partial \log f_j}{\partial \log N_j}\right)^{-1}\frac{\partial \log f_j}{\partial \tau_i}\right) \tag{B.39}$$

$$= \left(\sum_j \frac{N_j}{N}\left(\theta - \frac{\partial \log f_j}{\partial \log N_j}\right)^{-1}\right)^{-1}\left(\frac{N_i}{N}\left(\theta - \frac{\partial \log f_i}{\partial \log N_i}\right)^{-1}\frac{\partial \log f_i}{\partial \tau_i}\right), \tag{B.40}$$

where the last equality uses the fact that

$$\frac{\partial \log f_j}{\partial \tau_i} = 0, \quad \forall i \neq j. \tag{B.41}$$

Henceforth, we summarize this as

$$\frac{d\mathcal{U}}{d\tau_i} = \omega_i \frac{\partial \log f_i}{\partial \tau_i}, \quad \omega_i \equiv \frac{N_i\left(\theta - \frac{\partial \log f_i}{\partial \log N_i}\right)^{-1}}{\sum_j N_j\left(\theta - \frac{\partial \log f_j}{\partial \log N_j}\right)^{-1}}. \tag{B.42}$$

Now, on the right-hand side of the planner's first-order condition (B.33), these results imply

$$\sum_j \tau_j \frac{dN_j}{d\tau_i} = \left(\left(\theta - \frac{\partial \log f_i}{\partial \log N_i}\right)^{-1}\tau_i N_i - \omega_i \sum_j \left(\theta - \frac{\partial \log f_j}{\partial \log N_j}\right)^{-1}\tau_j N_j\right)\frac{\partial \log f_i}{\partial \tau_i}. \tag{B.43}$$

Substituting these results into the first-order condition, we have

$$\lambda N_i = \left(\omega_i\left(1 + \lambda \sum_j \left(\theta - \frac{\partial \log f_j}{\partial \log N_j}\right)^{-1}\tau_j N_j\right) - \lambda\left(\theta - \frac{\partial \log f_i}{\partial \log N_i}\right)^{-1}\tau_i N_i\right)\frac{\partial \log f_i}{\partial \tau_i}. \tag{B.44}$$

Rearranging to solve for τ_i, this becomes

$$\tau_i = -\left(\theta - \frac{\partial \log f_i}{\partial \log N_i}\right)\left(\frac{\partial \log f_i}{\partial \tau_i}\right)^{-1} + \frac{1 + \lambda \sum_j \left(\theta - \frac{\partial \log f_j}{\partial \log N_j}\right)^{-1} \tau_j N_j}{\lambda \sum_j \left(\theta - \frac{\partial \log f_j}{\partial \log N_j}\right)^{-1} N_j}. \tag{B.45}$$

Notice that the second term on the right-hand side is the same for all i. Thus, multiplying both sides by N_i, summing over i, and imposing budget balance implies

$$\tau_i = -\left(\theta - \frac{\partial \log f_i}{\partial \log N_i}\right)\left(\frac{\partial \log f_i}{\partial \tau_i}\right)^{-1} + \sum_j \frac{N_j}{N}\left(\theta - \frac{\partial \log f_j}{\partial \log N_j}\right)\left(\frac{\partial \log f_j}{\partial \tau_j}\right)^{-1}. \tag{B.46}$$

Locations receive a subsidy if their utility is relatively sensitive to transfers ($d \log f_i / d\tau_i$ is large) and relatively sensitive to population inflows ($d \log f_i / d \log N_i$ is higher).

To solve for transfers, we need only determine the derivatives in (B.46). Because $f_i(N_i, \tau_i) = x_i v(p_i)$ for some function $v(\cdot)$, we must have

$$\frac{\partial \log f_i}{\partial \log N_i} = \frac{\partial \log x_i}{\partial \log N_i} - \frac{p_i h_i}{x_i}\frac{\partial \log p_i}{\partial \log N_i} \tag{B.47}$$

$$\frac{\partial \log f_i}{\partial \tau_i} = \frac{\partial \log x_i}{\partial \tau_i} - \frac{p_i h_i}{x_i}\frac{\partial \log p_i}{\partial \tau_i}, \tag{B.48}$$

where both use the fact that housing expenditure shares are $p_i h_i / x_i = -\partial \log v(p_i) / \partial \log p_i$. To evaluate these, define the construction employment share $\hat{N}_{iX} = N_{iX}/N_i$ and note first that the construction first-order condition can be written

$$w_i \hat{N}_{iX} = p_i h_i \mu_i, \quad \mu_i \equiv \frac{\partial \log H_i(X_i, L_i)}{\partial \log X_i}, \tag{B.49}$$

hence $\pi_i = p_i h_i (1 - \mu_i)$ and so

$$x_i = w_i + p_i h_i (1 - \mu_i) + \tau_i. \tag{B.50}$$

Differentiating expenditures therefore gives (recall $w_i = A_i N_i^{\sigma}$)

$$\frac{\partial \log x_i}{\partial \log N_i} = \frac{1}{x_i}\left(w_i \sigma + p_i h_i (1 - \mu_i)\left(\frac{\partial \log p_i}{\partial \log N_i} + \frac{\partial \log h_i}{\partial \log N_i} + \frac{\partial \log(1 - \mu_i)}{\partial \log N_i}\right)\right). \tag{B.51}$$

Differentiating the construction first-order condition gives

$$\sigma + \frac{\partial \log \hat{N}_{iX}}{\partial \log N_i} = \frac{\partial \log p_i}{\partial \log N_i} + \frac{\partial \log h_i}{\partial \log N_i} + \frac{\partial \log \mu_i}{\partial \log N_i}, \tag{B.52}$$

and differentiating housing supply $h_i = H_i(A_X N_{iX}, L_i)/N_i$ gives

$$\frac{\partial \log h_i}{\partial \log N_i} = \mu_i \left(\frac{\partial \log \hat{N}_{iX}}{\partial \log N_i} + 1 \right) - 1, \tag{B.53}$$

therefore these jointly imply

$$\frac{\partial \log h_i}{\partial \log N_i} = \mu_i \left(\frac{\partial \log p_i}{\partial \log N_i} + \frac{\partial \log h_i}{\partial \log N_i} + \frac{\partial \log \mu_i}{\partial \log N_i} - \sigma \right) - (1 - \mu_i) \tag{B.54}$$

$$= \frac{\mu_i}{1 - \mu_i} \left(\frac{\partial \log p_i}{\partial \log N_i} + \frac{\partial \log \mu_i}{\partial \log N_i} - \sigma \right) - 1. \tag{B.55}$$

Substituting this back into the expenditure derivative, we then get

$$\frac{\partial \log x_i}{\partial \log N_i} = \frac{1}{x_i} \left(w_i \sigma + p_i h_i \left(\frac{\partial \log p_i}{\partial \log N_i} - \mu_i \sigma - (1 - \mu_i) \right) \right). \tag{B.56}$$

Then, using the fact that $Y_i/N_i = w_i \hat{N}_{iY} = w_i(1 - \hat{N}_{iX})$ and $\pi_i = p_i h_i (1 - \mu_i)$, we get that

$$\frac{\partial \log f_i}{\partial \log N_i} = \frac{1}{x_i} \left(\frac{Y_i}{N_i} \sigma - \pi_i \right). \tag{B.57}$$

Computing the derivative with respect to τ_i follows the same logic. We have the system

$$\frac{\partial \log x_i}{\partial \tau_i} = \frac{1}{x_i} \left(p_i h_i (1 - \mu_i) \left(\frac{\partial \log p_i}{\partial \tau_i} + \frac{\partial \log h_i}{\partial \tau_i} + \frac{\partial \log (1 - \mu_i)}{\partial \tau_i} \right) + 1 \right) \tag{B.58}$$

$$\frac{\partial \log \hat{N}_{iX}}{\partial \tau_i} = \frac{\partial \log p_i}{\partial \tau_i} + \frac{\partial \log h_i}{\partial \tau_i} + \frac{\partial \log \mu_i}{\partial \tau_i} \tag{B.59}$$

$$\frac{\partial \log h_i}{\partial \tau_i} = \mu_i \frac{\partial \log \hat{N}_{iX}}{\partial \tau_i}, \tag{B.60}$$

which together imply

$$\frac{\partial \log f_i}{\partial \tau_i} = \frac{1}{x_i}. \tag{B.61}$$

Now we can simply substitute this into the expression for the optimal transfer (B.46) to get

$$\tau_i = -\left(\theta - \frac{\partial \log f_i}{\partial \log N_i} \right) \left(\frac{\partial \log f_i}{\partial \tau_i} \right)^{-1} - \bar{\tau} \tag{B.62}$$

$$= -\theta x_i + \sigma \frac{Y_i}{N_i} - \pi_i - \bar{\tau}, \tag{B.63}$$

for budget-balancing constant $\bar{\tau}$. To get τ_i explicitly, substitute in $x_i = w_i + \pi_i + \tau_i$ to get

$$\tau_i = -\frac{\theta}{\theta + 1} w_i + \frac{\sigma}{\theta + 1} \frac{Y_i}{N_i} - \pi_i - \frac{\bar{\tau}}{1 + \theta}, \tag{B.64}$$

$$= -\frac{\theta}{\theta+1}(w_i - \bar{w}) + \frac{\sigma}{\theta+1}\left(\frac{Y_i}{N_i} - \frac{Y}{N}\right) - (\pi_i - \bar{\pi}), \tag{B.65}$$

where the second equality imposes budget balance.

Now, consider an alternative setting in which national profits $\bar{\pi}$ are rebated to households instead of π_i. We will use primes to denote this alternative allocation under a feasible transfer scheme $\{\tau_j'\}$. Conjecture that the optimal transfer in this setting takes the form

$$\tau_i' = -\frac{\theta}{\theta+1}(w_i' - \bar{w}') + \frac{\sigma}{\theta+1}\left(\frac{Y_i'}{N_i'} - \frac{Y'}{N}\right). \tag{B.66}$$

With this transfer policy and $w_i' = A_i(N_i')^\sigma$, the equilibrium $\{x_j', N_j', p_j', X_j'\}$ solves the system

$$x_i' = w_i' + \bar{\pi}' - \frac{\theta}{\theta+1}(w_i' - \bar{w}') + \frac{\sigma}{\theta+1}\left(\frac{Y_i'}{N_i'} - \frac{Y'}{N}\right) \tag{B.67}$$

$$\frac{N_i'}{N} = \frac{(x_i' v(p_i'))^{\frac{1}{\theta}}}{\sum_j (x_j' v(p_j'))^{\frac{1}{\theta}}} \tag{B.68}$$

$$w_i' = p_i' \frac{\partial H_i(X_i', L_i)}{\partial X_i'} A_X \tag{B.69}$$

$$p_i' H_i(X_i', L_i) = -x_i' N_i' \frac{\partial \log v(p_i')}{\partial \log p_i'}. \tag{B.70}$$

This system is identical to the system under τ_i when π_i is rebated, so these two policies in their respective settings must give rise to an identical equilibrium: $\{x_j', N_j', p_j', X_j'\} = \{x_j, N_j, p_j, X_j\}$. These policies therefore attain the same value of social welfare $\mathcal{U}$. Thus, because τ_i maximizes $\mathcal{U}$, and because the set of allocations feasible to the planner is the same in both settings, it must be that τ_i' maximizes $\mathcal{U}$.

Finally, consider what happens if we also had exogenous amenity spillovers so that utility becomes $U(c,h)N_i^{-\gamma}$. In this case, we now have $f_i(N_i, \tau_i) = x_i(N_i, \tau_i)v(p_i(N_i, \tau_i))N_i^{-\gamma}$ for indirect utility function v defined as above. The only part of the proof that then changes is that (B.47) becomes

$$\frac{\partial \log f_i}{\partial \log N_i} = \frac{\partial \log x_i}{\partial \log N_i} - \frac{p_i h_i}{x_i}\frac{\partial \log p_i}{\partial \log N_i} - \gamma, \tag{B.71}$$

and therefore the optimal transfer (B.46) becomes

$$\tau_i = -\left(\theta - \frac{\partial \log f_i}{\partial \log N_i}\right)\left(\frac{\partial \log f_i}{\partial \tau_i}\right)^{-1} - \bar{\tau} \tag{B.72}$$

$$= -(\theta+\gamma)x_i + \sigma\frac{Y_i}{N_i} - \pi_i - \bar{\tau} \tag{B.73}$$

$$= -\frac{\theta+\gamma}{\theta+\gamma+1}(w_i - \bar{w}) + \frac{\sigma}{\theta+\gamma+1}\left(\frac{Y_i}{N_i} - \frac{Y}{N}\right) - (\pi_i - \bar{\pi}). \tag{B.74}$$

Appendix C. Theory extensions

C.1 Robustness: Trade, commuting, and amenities

To assess the robustness of our scale-dependence results, we consider three extensions of the theory that are common in the spatial literature: trade, commuting, and amenities. In all three cases, we follow the conventional approaches to modeling these mechanisms and show that they do not introduce any new source of scale dependence.[3] They can change the level of equilibrium population shares, but they do not change the condition under which rescaling N leaves those shares unchanged.

Trade. Suppose that each location produces a differentiated variety, and that households consume these varieties through a CES aggregate: location i's consumption bundle is

$$c_i = \left(\sum_j c_{ij}^{\frac{\psi-1}{\psi}} \right)^{\frac{\psi}{\psi-1}}. \tag{C.1}$$

Let q_j denote the factory-gate price of the variety produced in j and let $\tau_{ij} \geq 1$ denote the iceberg cost of shipping that variety to i. The delivered price is $q_{ij} = \tau_{ij} q_j$, and the local goods price index is

$$Q_i = \left(\sum_j q_{ij}^{1-\psi} \right)^{\frac{1}{1-\psi}}. \tag{C.2}$$

The household problem is then exactly the problem from Section 2.3, except that the price of the consumption composite is Q_i rather than one. Defining the real wage and relative house price by

$$\tilde{w}_i \equiv \frac{w_i}{Q_i}, \qquad \tilde{p}_i \equiv \frac{p_i}{Q_i}, \tag{C.3}$$

indirect utility is

$$U_i^* = \tilde{w}_i v(\tilde{p}_i), \tag{C.4}$$

and population shares solve

$$\frac{N_i}{N} = \frac{(\tilde{w}_i v(\tilde{p}_i))^{\frac{1}{\theta}}}{\sum_k (\tilde{w}_k v(\tilde{p}_k))^{\frac{1}{\theta}}}. \tag{C.5}$$

Thus, on the household side, trade only replaces nominal wages and house prices with real wages and relative house prices.

[3] Of course, this does not mean that there are not other modeling approaches that would lead to scale dependence. Indeed, we consider two such approaches in Appendices C.3 and C.4 below.

On the production side, each variety is produced as (11) with perfect competition, so

$$w_j = q_j A_j N_j^{\sigma}. \tag{C.6}$$

CES demand implies that the share of location i's goods expenditure spent on variety j is

$$\lambda_{ij} = \left(\frac{q_{ij}}{Q_i}\right)^{1-\psi}. \tag{C.7}$$

Goods-market clearing therefore takes the form

$$q_j A_j N_j^{\sigma} N_{jY} = \sum_i \lambda_{ij} E_i, \tag{C.8}$$

where E_i is total expenditure on the differentiated goods by households in i.[4]

Now define $s_i \equiv N_i/N$, $s_{iY} \equiv N_{iY}/N$, and $e_i \equiv E_i/(N_i \tilde{w}_i)$ as the goods expenditure share. If all population shares are held fixed and N is rescaled, then (C.6)–(C.8) are homogeneous in exactly the same way as the baseline goods-market condition: wages scale with N^{σ}, while the relative goods prices $\{q_{ij}/Q_i\}_{i,j}$ are functions of population shares, employment shares, and the goods expenditure shares. Hence real wages can be written

$$\tilde{w}_i = \omega_i(\{s_j, s_{jY}, e_j\}_j) N^{\sigma}, \tag{C.9}$$

for some functions $\omega_i(\cdot)$ that do not depend directly on N.

This is the key point. The Armington block changes the endogenous term multiplying N^{σ}, but it does not add a new power of N. The housing block is also unchanged after replacing $\{w_i, p_i\}$ by $\{\tilde{w}_i, \tilde{p}_i\}$: demand is

$$h_i^{\mathrm{D}}(\tilde{w}_i, \tilde{p}_i) = -\frac{\tilde{w}_i}{\tilde{p}_i} \frac{\partial \log v(\tilde{p}_i)}{\partial \log \tilde{p}_i}, \tag{C.10}$$

and the construction first-order condition is

$$\tilde{w}_i = \tilde{p}_i H_i'(X_i) A_X. \tag{C.11}$$

Therefore, after normalizing by the local goods price index, the equilibrium system is the same as in Section 2.3, with A_i replaced by an endogenous market-access term $\omega_i(\cdot)$ that depends on population and expenditure shares. Since this term carries no explicit dependence on N, Armington trade does

[4] As in the main text, absentee land rents leave the economy in the form of lost goods. Including those lost goods on the right-hand side of (C.8) would not affect the argument, since the lost-goods term is also homogeneous of degree one in nominal income and rents.

not change the scale-independence conditions in Proposition 3. In particular, Cobb-Douglas utility and housing supply still imply scale independence, while deviations from Cobb-Douglas demand or supply generate scale dependence for the same reason as in the baseline model.

Commuting. Now suppose that households can live in i and work in j. Commuting from i to $j \neq i$ results in a proportional utility cost $\tau_{ij} > 1$, so that the household solves

$$V_k = \max_{i,j,c_{ijk},h_{ijk}} \frac{U(c_{ijk}, h_{ijk})\epsilon_{ijk}}{\tau_{ij}}, \tag{C.12}$$

$$w_j \geq c_{ijk} + p_i h_{ijk}. \tag{C.13}$$

Let N_{ij} denote the number households living in i and working in j, implying residence and workplace populations

$$\bar{N}_i \equiv \sum_j N_{ij}, \qquad \underline{N}_j \equiv \sum_i N_{ij}. \tag{C.14}$$

Pair-specific indirect utility is then

$$U^*_{ij} = \frac{w_j v(p_i)}{\tau_{ij}}. \tag{C.15}$$

The pair shares therefore solve

$$\frac{N_{ij}}{N} = \frac{\left(w_j v(p_i)\tau_{ij}^{-1}\right)^{\frac{1}{\theta}}}{\sum_{k,l}\left(w_l v(p_k)\tau_{kl}^{-1}\right)^{\frac{1}{\theta}}}. \tag{C.16}$$

The only change relative to the baseline is that the wage relevant for income is the workplace wage, while the house price relevant for housing demand is the residence price.

Let $\bar{s}_i \equiv \bar{N}_i/N$, $\underline{s}_j \equiv \underline{N}_j/N$, and $n_{ij} \equiv N_{ij}/N$. Goods production in workplace j gives

$$w_j = A_j \underline{N}_j^{\sigma} = N^{\sigma} A_j \underline{s}_j^{\sigma}. \tag{C.17}$$

Since preferences are homogeneous, housing demand is linear in income. Hence, aggregate housing demand in residence i is

$$H_i(X_i, L_i) = \sum_j N_{ij} h_{ij}^{\mathrm{D}} = -\frac{\bar{N}_i \bar{w}_i}{p_i} \frac{\partial \log v(p_i)}{\partial \log p_i}, \tag{C.18}$$

where

$$\bar{w}_i \equiv \sum_j \frac{N_{ij}}{\bar{N}_i} w_j \tag{C.19}$$

is average labor income among residents of i. By (C.17), $\bar{w}_i = \bar{\omega}_i(\{n_{ij}\}_{i,j})N^{\sigma}$ for some function $\bar{\omega}_i(\cdot)$ that depends only on pair shares.

Equations (C.16)–(C.18) show that commuting adds only share-dependent terms, like in the Armington trade model. To see the cancellation most transparently, suppose that utility and housing supply take the Cobb-Douglas forms from Proposition 3. Then $p_i h_{ij} = \alpha w_j$, and combining housing-market clearing with the construction first-order condition gives

$$X_i = \alpha(1-\nu) A_X \frac{\bar{w}_i}{w_i} \bar{N}_i. \tag{C.20}$$

Thus structures in each residence are proportional to aggregate population times a function of pair shares. It follows that p_i, $v(p_i)$, and $w_j v(p_i)$ can all be written as common powers of N times functions of $\{n_{ij}\}_{i,j}$. The common powers cancel from (C.16), leaving an equilibrium system entirely in pair shares. Hence the Cobb-Douglas benchmark remains scale-independent.

With non-Cobb-Douglas demand or supply, the same objects that generated scale dependence in Proposition 3 reappear in (C.18): housing expenditure shares and local house-price responses. Commuting costs alter the average income $\bar{w}_i$ and the mapping between residence and workplace populations, but these are functions of shares and bilateral costs, not independent powers of aggregate population. Therefore commuting can change the level of concentration, but it does not change the economic source of scale dependence.

Amenities. Finally, suppose that local amenities multiply utility and take the standard power form $a_i N_i^\gamma$, where a_i and γ are constants.[5] Indirect utility in location i is then

$$a_i N_i^\gamma U^*(w_i, p_i) = a_i N_i^\gamma w_i v(p_i). \tag{C.21}$$

After substituting $w_i = A_i N_i^\sigma$, population shares satisfy

$$\frac{N_i}{N} = \frac{(a_i A_i v(p_i))^{\frac{1}{\theta-\sigma-\gamma}}}{\sum_j (a_j A_j v(p_j))^{\frac{1}{\theta-\sigma-\gamma}}}, \tag{C.22}$$

provided $\theta > \sigma + \gamma$. Thus a power amenity simply changes the effective strength of the location-specific scale term, replacing $\theta - \sigma$ with $\theta - \sigma - \gamma$. Since γ is common across locations, N_i^γ is another common power of aggregate population times a function of shares. It can affect the level of concentration, but it does not create scale dependence unless the amenity itself departs from this common power form.

[5] These amenity spillovers are often taken to be exogenous, but they can also arise endogenously from local spending on public goods. Appendix C.4 below derives such a result.

C.2 Non-homothetic utility

Scale dependence. Assume the setting from Section 2.3, but relax the assumption that $U(c_i, h_i)$ is homothetic. Now, we can no longer use the separation $U^*(w_i, p_i) = w_i v(p_i)$. Define the income and house-price elasticities of utility[6]

$$\gamma_{wi} \equiv \frac{U_w^*(w_i, p_i) w_i}{U^*(w_i, p_i)} \quad \text{and} \quad \gamma_{pi} \equiv -\frac{U_p^*(w_i, p_i) p_i}{U^*(w_i, p_i)}. \tag{C.23}$$

By Roy's identity, we have housing demand

$$h_i^{\mathrm{D}}(w_i, p_i) = -\frac{U_p^*(w_i, p_i)}{U_w^*(w_i, p_i)} = \frac{w_i}{p_i} \frac{\gamma_{pi}}{\gamma_{wi}}. \tag{C.24}$$

Differentiating the population share $N_i/N \propto U^*(w_i, p_i)^{1/\theta}$ then gives

$$\frac{d \log N_i}{d \log N} = \frac{1}{\theta} \left(\gamma_{wi} \frac{d \log w_i}{d \log N_i} - \gamma_{pi} \frac{d \log p_i}{d \log N_i} \right) \frac{d \log N_i}{d \log N} + \tilde{\Gamma} \tag{C.25}$$

$$= \left(\theta - \gamma_{wi} \left(\sigma - \frac{p_i h_i}{w_i} \frac{d \log p_i}{d \log N_i} \right) \right)^{-1} \Gamma. \tag{C.26}$$

The rest follows the proof of Proposition 3. The total housing demand change is

$$\frac{d \log h_i^{\mathrm{D}}}{d \log N_i} = -\zeta_i^{\mathrm{D}} \frac{d \log p_i}{d \log N_i} + \xi_i \sigma \tag{C.27}$$

where

$$\xi_i \equiv \frac{\partial \log h_i^{\mathrm{D}}}{\partial \log w_i} = 1 + \frac{\partial \log \gamma_{pi}}{\partial \log w_i} - \frac{\partial \log \gamma_{wi}}{\partial \log w_i}. \tag{C.28}$$

The housing supply change is, as before, (B.12). Equating supply and demand then gives

$$\frac{d \log p_i}{d \log N_i} = \frac{1}{\zeta_i^{\mathrm{S}} + \zeta_i^{\mathrm{D}}} + \frac{\zeta_i^{\mathrm{S}} + \xi_i}{\zeta_i^{\mathrm{S}} + \zeta_i^{\mathrm{D}}} \sigma. \tag{C.29}$$

Putting all of this together, the total population change is

$$\frac{d \log N_i}{d \log N} = \left(\theta - \gamma_{wi} \left(\sigma - \frac{p_i h_i}{w_i} \left(\frac{1}{\zeta_i^{\mathrm{S}} + \zeta_i^{\mathrm{D}}} + \frac{\zeta_i^{\mathrm{S}} + \xi_i}{\zeta_i^{\mathrm{S}} + \zeta_i^{\mathrm{D}}} \sigma \right) \right) \right)^{-1} \Gamma. \tag{C.30}$$

Quantitative importance. (C.30) shows that with non-homothetic utility, heterogeneity in income elasticities of utility (γ_{wi}) and housing demand (ξ_i) can generate scale dependence, even with constant $\{p_i h_i / w_i, \zeta_i^{\mathrm{S}}, \zeta_i^{\mathrm{D}}\}$. However, the scale dependence created by γ_{wi} and ξ_i is likely small for

[6] The income elasticity γ_{wi} also equals $(U_c(c_i, h_i) c_i + U_h(c_i, h_i) h_i)/U(c_i, h_i)$.

the following reasons. Focusing first on the role of γ_{wi}, it multiplies two terms (C.30). The first term, $\gamma_{wi}\sigma$, is small relative to the other terms in (C.30) because σ is small. The role of the second term depends on how γ_{wi} interacts with housing expenditure shares. To determine its importance, we can use the standard formulation of non-homothetic housing demand, non-homothetic PIGL (Robert-Nicoud et al. 2026), for which indirect utility is

$$U^*(w_i, p_i) = \exp\left\{\frac{w_i^{1-\xi_w} - 1}{1-\xi_w} - \alpha\frac{p_i^{1-\zeta_p} - 1}{1-\zeta_p}\right\}. \tag{C.31}$$

These preferences imply $\gamma_{wi} = w_i^{1-\xi_w}$, $\gamma_{pi} = \alpha p_i^{1-\zeta_p}$, $h_i^{\mathrm{D}}(w_i, p_i) = \alpha w_i^{\xi_w} p_i^{-\zeta_p}$, and therefore $\gamma_{wi}\frac{p_i h_i}{w_i} = \alpha p_i^{1-\zeta_p}$. This implies that the scale dependence arises under the exact same conditions as in the homothetic PIGL case—when $p_i^{1-\zeta_p}$ varies with N_i—because the γ_{wi} shows up in the denominator of the expenditure share and therefore drops out of $\gamma_{wi}\frac{p_i h_i}{w_i}$. Finally, turning to the role of ξ_i, this is likely to have a small effect given that standard estimates of the income elasticity of housing demand are between 0.6 and 0.8 (Albouy et al. 2016; Combes et al. 2019; Finlay and Williams 2025), close to estimates of price elasticities (Robert-Nicoud et al. 2026), and any effect is scaled by σ, which is small.

C.3 Scale dependence from multiple goods

To study the consequences of multiple goods for spatial concentration, we embed a simplified version of the model of structural transformation and urbanization from Michaels et al. (2012) into our static model.

Assumptions. The setting is as in Section 2.2. To isolate the force in the cleanest possible way, we abstract from housing ($\alpha = 0$). We replace the single traded good c_i with two traded goods c_{is}, agriculture ($s = a$) and non-agriculture ($s = n$). Households consume a CES composite

$$c_i = \left((1-\alpha_c)c_{in}^{\frac{\eta_c-1}{\eta_c}} + \alpha_c c_{ia}^{\frac{\eta_c-1}{\eta_c}}\right)^{\frac{\eta_c}{\eta_c-1}}, \tag{C.32}$$

with $\eta_c < 1$ corresponding to the standard structural-transformation case in which agricultural and non-agricultural goods are complements. Goods are freely traded, so their prices $\{p_a, p_n\}$ are common across locations. Conditional on income w_{is}, households solve

$$w_{is} \geq p_n c_{in} + p_a c_{ia}. \tag{C.33}$$

The associated price index $P(p_a, p_n)$ is normalized to one. As in the main text, households receive i.i.d. Fréchet location shocks with dispersion θ.

To simplify the math, we define locations to be arbitrarily small plots of land of equal measure $L_i = L$, so that there is a continuum of locations $i \in [0, J]$. We order these locations by their innate productivity, A_i, which is continuously differentiable and strictly increasing in i. While i is technically just an indexing convention, it is intuitive to think about A_i as being positively spatially autocorrelated, so that two adjacent plots of land are likely to have similar A_i.[7] Because plots of land are small, we assume that they each specialize in only one sector $s \in \{a, n\}$. If location i specializes in sector s, total output in that location will be[8]

$$Y_{is} = A_i \Gamma_s N_i^{1-\mu_s} L^{\mu_s}, \quad 0 < \mu_n < \mu_a < 1. \tag{C.34}$$

Non-agriculture is less land-intensive than agriculture, so it has stronger returns to population density. We could also allow for sector-specific agglomeration elasticities, replacing $A_i \Gamma_s$ with $A_i \Gamma_s N_i^{\sigma_s}$ and assuming $\sigma_n > \sigma_a$, as in Michaels et al. (2012). The effect of $\sigma_n > \sigma_a$ is qualitatively the same as the effect of $\mu_n < \mu_a$, since both increase returns to density.

Firms are competitive and land rents accrue to absentee landowners. The wage paid in a location specializing in s thus equals the marginal product of labor,

$$w_{is} = (1 - \mu_s) p_s \Gamma_s A_i \left(\frac{N_i(s_i)}{L} \right)^{-\mu_s}. \tag{C.35}$$

The sectoral choice of a location is therefore

$$s_i \in \underset{s \in \{a,n\}}{\arg\max}\, w_{is}. \tag{C.36}$$

Note that, in equilibrium, $N_i(s_i)$ will depend on which sector is chosen in i, and this is internalized in the specialization choice.

Equilibrium. Given sector choices s_i, population allocations solve

$$\frac{N_i}{N} = \frac{w_{is_i}^{1/\theta}}{\int_0^J w_{js_j}^{1/\theta} dj}. \tag{C.37}$$

Combining (C.35)–(C.36) gives a simple density cutoff: defining

$$\bar{n} \equiv \left(\frac{(1 - \mu_a) p_a \Gamma_a}{(1 - \mu_n) p_n \Gamma_n} \right)^{\frac{1}{\mu_a - \mu_n}}, \tag{C.38}$$

[7] One can define larger regions (e.g., prefectures or MSAs) as collections of adjacent land plots. With positive spatial autocorrelation, these larger regions would have heterogeneous sectoral mixes and populations.

[8] Notice that each location has the same relative productivity in agriculture and non-agriculture Γ_a / Γ_n, but differs in its overall productivity A_i. This simplifies the results but is not necessary.

location i specializes in non-agriculture if and only if

$$\frac{N_i}{L} \geq \bar{n}. \tag{C.39}$$

Thus, in any equilibrium, all non-agricultural regions will have higher density than all agricultural regions. Let $\bar{\imath} \equiv \{i \in [0, J] : N_{\bar{\imath}}/L = \bar{n}\} \in (0, J)$.[9]

The goods prices are determined by national goods-market clearing. The first-order conditions for (C.32) imply

$$\frac{p_a}{p_n} = \frac{\alpha_c}{1-\alpha_c}\left(\frac{Y_n}{Y_a}\right)^{\frac{1}{\eta_c}}, \tag{C.40}$$

Using the sorting rule (C.39) and the fact that it maps one-to-one with productivity A_i, this can be written as

$$\frac{p_a}{p_n} = \frac{\alpha_c}{1-\alpha_c}\left(\frac{\Gamma_n L^{\mu_n}}{\Gamma_a L^{\mu_a}}\right)^{\frac{1}{\eta_c}}\left(\frac{\int_{\bar{\imath}}^{J} A_i N_i^{1-\mu_n} di}{\int_0^{\bar{\imath}} A_i N_i^{1-\mu_a} di}\right)^{\frac{1}{\eta_c}}. \tag{C.41}$$

Substituting this back into the wages (C.35) and density threshold (C.39) with the conditions (C.37) and $N_{\bar{\imath}}/L = \bar{n}$ then determines the equilibrium.

Scale dependence. We first study how the cutoff $\bar{\imath}(N) \in (0, J)$ changes with aggregate population N. Since A_i is strictly increasing, locations $i < \bar{\imath}(N)$ specialize in agriculture and locations $i > \bar{\imath}(N)$ specialize in non-agriculture. The cutoff location has density $\bar{n}(N)$. Using (C.35) and (C.37), equilibrium densities can be written as

$$\frac{N_i}{L} = \begin{cases} \bar{n}(N)\left(\dfrac{A_i}{A_{\bar{\imath}(N)}}\right)^{\frac{1}{\theta+\mu_a}} & i < \bar{\imath}(N), \\ \bar{n}(N)\left(\dfrac{A_i}{A_{\bar{\imath}(N)}}\right)^{\frac{1}{\theta+\mu_n}} & i > \bar{\imath}(N). \end{cases} \tag{C.42}$$

The goods-market clearing condition (C.40), the specialization cutoff (C.38), and the adding-up condition $\int_0^J N_i di = N$ jointly determine $\{\bar{\imath}(N), \bar{n}(N), p_a(N)/p_n(N)\}$. Differentiating (C.42), we have

$$\frac{d\log N_i}{d\log N} = \begin{cases} \frac{\partial \log \bar{n}(N)}{\partial \log N} - \frac{1}{\theta+\mu_a}\frac{\partial \log A_{\bar{\imath}(N)}}{\partial \log N} & i < \bar{\imath}(N), \\ \frac{\partial \log \bar{n}(N)}{\partial \log N} - \frac{1}{\theta+\mu_n}\frac{\partial \log A_{\bar{\imath}(N)}}{\partial \log N} & i > \bar{\imath}(N). \end{cases} \tag{C.43}$$

Hence, the direction of scale dependence depends entirely on the sign of the response of cutoffs $\bar{\imath}(N)$ and $\bar{n}(N)$. The following proposition shows that, when $\eta_c \in (0, 1)$ and $\mu_a > \mu_n$, both $\bar{\imath}(N)$ and

[9] The equilibrium will always have an interior solution $\bar{\imath} \in (0, J)$, because, under (C.32), zero aggregate supply of either good would imply an infinite relative price for that good. Since every location can produce either good, such a corner cannot be an equilibrium, so any equilibrium must feature $\bar{\imath}(N) \in (0, J)$.

$\bar{n}(N)$ are increasing in N, and depopulation leads to rising spatial concentration.

Proposition 6 (Endogenous specialization and spatial concentration)**.** *Agricultural regions are strictly less dense than non-agricultural regions. Moreover, if $\eta_c \in (0,1)$ and $\mu_a > \mu_n$, then in response to a decline in total population N: (i) the cutoff $\bar{\imath}(N)$ falls, so a positive measure of marginal locations switches from agriculture to non-agriculture; (ii) population shares rise in continuing non-agricultural regions and fall in continuing agricultural regions; and, as a result, (iii) spatial concentration rises.*

Proof. The first statement follows directly from (C.39): all agricultural regions are below the density cutoff and all non-agricultural regions are above it.

The population changes (C.43) tell us that all of the claims (i)–(iii) of the proposition depend on the signs of $\bar{\imath}'(N)$ and $\bar{n}'(N)$. We therefore establish the following intermediate result, then show that the statements of the proposition follow from it.

Lemma 2 (Cutoff comparative statics)**.** *Suppose $\eta_c \in (0,1)$ and $\mu_a > \mu_n$. Then $\bar{\imath}'(N) > 0$ and $\bar{n}'(N) > 0$. Moreover, $d(N_i/N)/dN > 0$ for every continuing agricultural location ($i < \bar{\imath}(N)$) and $d(N_i/N)/dN < 0$ for every continuing non-agricultural location ($i > \bar{\imath}(N)$).*

Proof. First, use (C.42) to write aggregate population as

$$N = \left(\int_0^{\bar{\imath}} \left(\frac{A_i}{A_{\bar{\imath}}} \right)^{\frac{1}{\theta+\mu_a}} di + \int_{\bar{\imath}}^{J} \left(\frac{A_i}{A_{\bar{\imath}}} \right)^{\frac{1}{\theta+\mu_n}} di \right) L\bar{n}. \tag{C.44}$$

Next, substitute (C.42) into sectoral output:

$$Y_a = L\Gamma_a \bar{n}^{1-\mu_a} \int_0^{\bar{\imath}} A_i \left(\frac{A_i}{A_{\bar{\imath}}} \right)^{\frac{1-\mu_a}{\theta+\mu_a}} di, \tag{C.45}$$

$$Y_n = L\Gamma_n \bar{n}^{1-\mu_n} \int_{\bar{\imath}}^{J} A_i \left(\frac{A_i}{A_{\bar{\imath}}} \right)^{\frac{1-\mu_n}{\theta+\mu_n}} di. \tag{C.46}$$

Combining these expressions with goods-market clearing (C.40) and the cutoff equation (C.38) gives

$$\bar{n}^{(\mu_a-\mu_n)\frac{\eta_c-1}{\eta_c}} = \Omega_c \left(\frac{\int_{\bar{\imath}}^{J} A_i^{\frac{\theta+1}{\theta+\mu_n}} di}{\int_0^{\bar{\imath}} A_i^{\frac{\theta+1}{\theta+\mu_a}} di} \right)^{\frac{1}{\eta_c}} A_{\bar{\imath}}^{-\left(\frac{1-\mu_n}{\theta+\mu_n} - \frac{1-\mu_a}{\theta+\mu_a}\right)\frac{1}{\eta_c}}, \tag{C.47}$$

where $\Omega_c > 0$ collects constants. The ratio on the right-hand side is strictly decreasing in $\bar{\imath}$, both because the integral in the numerator falls and the integral in the denominator rises, and because $A_{\bar{\imath}}$ rises and its exponent is negative (since $\mu_a > \mu_n$). Because $\eta_c < 1$, the exponent on the left-hand side of (C.47) is negative. Hence $\bar{n}$ is strictly increasing in $\bar{\imath}$.

We first show that $\bar{\imath}$ is strictly increasing in N. Equation (C.47) determines the threshold density $\bar{n}$ associated with that cutoff, and equation (C.44) then gives the aggregate population level consistent with the cutoff. We will show that this population level is strictly increasing in the candidate cutoff, and then, since the equilibrium cutoff $\bar{\imath}(N)$ is the inverse of this relationship, this implies $\bar{\imath}'(N) > 0$. Substitute $\bar{n}(\bar{\imath})$ from (C.47) into (C.44). Differentiating (C.47) gives

$$\frac{d\log\bar{n}}{d\bar{\imath}} = \frac{A_{\bar{\imath}}^{\frac{\theta+1}{\theta+\mu_n}}\left(\int_{\bar{\imath}}^{J} A_i^{\frac{\theta+1}{\theta+\mu_n}}di\right)^{-1} + A_{\bar{\imath}}^{\frac{\theta+1}{\theta+\mu_a}}\left(\int_{0}^{\bar{\imath}} A_i^{\frac{\theta+1}{\theta+\mu_a}}di\right)^{-1}}{(\mu_a-\mu_n)(1-\eta_c)} + \frac{\left(\frac{1-\mu_n}{\theta+\mu_n}-\frac{1-\mu_a}{\theta+\mu_a}\right)\frac{d\log A_{\bar{\imath}}}{d\bar{\imath}}}{(\mu_a-\mu_n)(1-\eta_c)}, \tag{C.48}$$

which is strictly positive. Moreover, since

$$\frac{\frac{1-\mu_n}{\theta+\mu_n}-\frac{1-\mu_a}{\theta+\mu_a}}{\mu_a-\mu_n} = \frac{\theta+1}{(\theta+\mu_n)(\theta+\mu_a)} > \frac{1}{\theta+\mu_n}, \tag{C.49}$$

equation (C.48) implies

$$\frac{d\log\bar{n}}{d\bar{\imath}} > \frac{1}{\theta+\mu_n}\frac{d\log A_{\bar{\imath}}}{d\bar{\imath}}. \tag{C.50}$$

Next, differentiating (C.44) gives[10]

$$\frac{d\log N}{d\bar{\imath}} = \frac{d\log\bar{n}}{d\bar{\imath}} - \frac{d\log A_{\bar{\imath}}}{d\bar{\imath}}\frac{\frac{1}{\theta+\mu_a}\int_0^{\bar{\imath}}\left(\frac{A_i}{A_{\bar{\imath}}}\right)^{\frac{1}{\theta+\mu_a}}di + \frac{1}{\theta+\mu_n}\int_{\bar{\imath}}^{J}\left(\frac{A_i}{A_{\bar{\imath}}}\right)^{\frac{1}{\theta+\mu_n}}di}{\int_0^{\bar{\imath}}\left(\frac{A_i}{A_{\bar{\imath}}}\right)^{\frac{1}{\theta+\mu_a}}di + \int_{\bar{\imath}}^{J}\left(\frac{A_i}{A_{\bar{\imath}}}\right)^{\frac{1}{\theta+\mu_n}}di}. \tag{C.51}$$

The fraction multiplying $d\log A_{\bar{\imath}}/d\bar{\imath}$ in (C.51) is a weighted average of $1/(\theta+\mu_a)$ and $1/(\theta+\mu_n)$, and is therefore strictly less than $1/(\theta+\mu_n)$. Combining this with (C.50) gives

$$\frac{d\log N}{d\bar{\imath}} > 0. \tag{C.52}$$

Hence $\bar{\imath}'(N) > 0$. Since $\bar{n}$ is strictly increasing in $\bar{\imath}$, it also follows that $\bar{n}'(N) > 0$.

Finally, consider a continuing location. From (C.42),

$$\frac{d\log(N_i/L)}{d\bar{\imath}} = \frac{d\log\bar{n}}{d\bar{\imath}} - \frac{1}{\theta+\mu_s}\frac{d\log A_{\bar{\imath}}}{d\bar{\imath}}, \qquad s\in\{a,n\}. \tag{C.53}$$

By (C.50), this expression is strictly positive for both sectors. Therefore continuing locations become denser when N rises and less dense when N falls.

[10]The boundary terms cancel because the cutoff location has density $\bar{n}$ whether it is assigned to agriculture or non-agriculture.

To study population shares, subtract (C.51) from (C.53). Since L is fixed,

$$\frac{d\log(N_i/N)}{d\bar{\imath}} = \left[\frac{\frac{1}{\theta+\mu_a}\int_0^{\bar{\imath}}\left(\frac{A_i}{A_{\bar{\imath}}}\right)^{\frac{1}{\theta+\mu_a}}di + \frac{1}{\theta+\mu_n}\int_{\bar{\imath}}^{J}\left(\frac{A_i}{A_{\bar{\imath}}}\right)^{\frac{1}{\theta+\mu_n}}di}{\int_0^{\bar{\imath}}\left(\frac{A_i}{A_{\bar{\imath}}}\right)^{\frac{1}{\theta+\mu_a}}di + \int_{\bar{\imath}}^{J}\left(\frac{A_i}{A_{\bar{\imath}}}\right)^{\frac{1}{\theta+\mu_n}}di} - \frac{1}{\theta+\mu_s}\right]\frac{d\log A_{\bar{\imath}}}{d\bar{\imath}}. \tag{C.54}$$

The term in brackets is positive for agricultural locations, since $1/(\theta+\mu_a)$ is below the weighted average. It is negative for non-agricultural locations, since $1/(\theta+\mu_n)$ is above the weighted average. Since $d\log N/d\bar{\imath} > 0$, the same signs apply to $d\log(N_i/N)/dN$ and hence to $d(N_i/N)/dN$. Thus $d(N_i/N)/dN > 0$ for every continuing agricultural location and $d(N_i/N)/dN < 0$ for every continuing non-agricultural location. □

Statements (i) and (ii) follow from Lemma 2: if $dN < 0$, then $d\bar{\imath} < 0$ and relative population rises in continuing non-agricultural regions and falls in continuing agricultural regions. Since A_i is continuous and strictly increasing, the interval $(\bar{\imath}(N+dN), \bar{\imath}(N))$ has positive measure for any $dN < 0$ and consists of locations that switch from agriculture to non-agriculture. By the same reasoning as in the proof of Lemma 1, statement (iii) then follows from the fact that population shares rise in the high-density locations and fall in the low-density locations. □

Agriculture is intensive in land and not people, and non-agriculture is intensive in people and not land. It therefore makes sense that denser places specialize in the latter. When aggregate population falls, holding technology fixed, the economy can economize on land-intensive production by converting marginal agricultural locations into non-agricultural locations and concentrating population shares in high-productivity, high-density locations. Under that reallocation, the relative population in non-agricultural places rises and so total spatial concentration rises.

Importantly, the analysis above assumes no structural change—that is, no change in technologies $\{A_a, A_n\}$ or in preferences $\{\alpha_c\}$. We believe this is a reasonable description of developed economies in the 21st century, since these economies already had very low employment shares in agriculture. Likewise, we believe this is a reasonable assumption for the predicted rapid transition to a world with lower population, since it seems unlikely that a world with a collapsing population should see a high rate of technological innovation (Jones 2022; Peters and Walsh 2026).

C.4 Scale dependence from multiple amenities

To study scale effects in amenities, we consider a simple extension in which local governments provide two types of amenities. The first type is *divisible*: the service flow to each resident is proportional to per-capita spending. The second type is *shared-capacity*: the service flow from a given amount of per-capita spending is higher in larger places because fixed facilities, equipment,

or coverage networks can be used by many residents before congestion fully binds. Put differently, shared-capacity amenities are less congestible, meaning that their cost per resident falls with the number of users.

To fix ideas, first consider a few examples. Hospital networks and transit systems are shared-capacity amenities, but they must be paired with divisible inputs like day-to-day staffing and maintenance. School buildings and library branches are shared-capacity amenities, but they must be paired with teachers, books, and programs. Emergency-response coverage is a shared-capacity amenity, but it is effective only when paired with enough personnel. A common theme among these examples is that the divisible and shared-capacity amenities are complements, a feature we will show naturally generates scale dependence.

Assumptions. To isolate the amenity force, abstract from housing ($\alpha = 0$). A local government in each location taxes labor income at rates τ_{di} and τ_{si} to fund divisible amenities (d) and shared-capacity amenities (s), respectively. Total taxes are $\tau_i \equiv \tau_{di} + \tau_{si} < 1$, so private consumption is

$$c_i = (1 - \tau_i) w_i. \tag{C.55}$$

Divisible amenity services equal per-capita spending, $\tau_{di} w_i$. Shared-capacity amenity services equal per-capita spending times a scale term, $\tau_{si} w_i N_i^{\chi}$, where $\chi \in (0, 1)$ captures partial sharing. Holding wages fixed, the per-resident spending needed to deliver a given service level falls with local population at elasticity χ. The limiting case $\chi = 1$ would mean a pure fixed-capacity amenity over the relevant range: doubling the number of residents does not require higher total spending to maintain the same per-resident service flow. We assume $\chi < 1$, so shared-capacity amenities exhibit economies of scale but are still congestible.

Households derive utility from the two amenity services through a CES composite

$$g_i = w_i \left((1 - \nu_g) \tau_{di}^{\frac{\eta_g - 1}{\eta_g}} + \nu_g (\tau_{si} N_i^{\chi})^{\frac{\eta_g - 1}{\eta_g}} \right)^{\frac{\eta_g}{\eta_g - 1}}, \tag{C.56}$$

where $\nu_g \in (0, 1)$. The case $\eta_g < 1$ means that divisible and shared-capacity amenities are complements. In this case, losing population makes it harder to sustain the scalable facility or network, and this is especially costly when the more divisible inputs cannot easily substitute for it. Household utility before the location taste shock is

$$c_i^{1 - \alpha_g} g_i^{\alpha_g}, \qquad \alpha_g \in (0, 1). \tag{C.57}$$

As in the rest of the paper, wages are $w_i = A_i N_i^{\sigma}$.

Equilibrium. The local government chooses $\{\tau_{di}, \tau_{si}\}$ to maximize household utility, taking N_i and w_i as given. Since w_i enters both c_i and g_i multiplicatively, the problem reduces to

$$\max_{\tau_{di},\tau_{si}\geq 0} \left\{ (1-\tau_{di}-\tau_{si})^{1-\alpha_g} \left((1-\nu_g)\tau_{di}^{\frac{\eta_g-1}{\eta_g}} + \nu_g(\tau_{si}N_i^{\chi})^{\frac{\eta_g-1}{\eta_g}} \right)^{\frac{\alpha_g\eta_g}{\eta_g-1}} \right\}. \tag{C.58}$$

Let $\{\tau_{di}^*, \tau_{si}^*\}$ denote the solution. The first-order conditions imply

$$(1-\alpha_g)\frac{\tau_{di}^*}{1-\tau_{di}^*-\tau_{si}^*} = \alpha_g \frac{(1-\nu_g)\tau_{di}^{*\frac{\eta_g-1}{\eta_g}}}{(1-\nu_g)\tau_{di}^{*\frac{\eta_g-1}{\eta_g}} + \nu_g(\tau_{si}^*N_i^{\chi})^{\frac{\eta_g-1}{\eta_g}}}, \tag{C.59}$$

$$(1-\alpha_g)\frac{\tau_{si}^*}{1-\tau_{di}^*-\tau_{si}^*} = \alpha_g \frac{\nu_g(\tau_{si}^*N_i^{\chi})^{\frac{\eta_g-1}{\eta_g}}}{(1-\nu_g)\tau_{di}^{*\frac{\eta_g-1}{\eta_g}} + \nu_g(\tau_{si}^*N_i^{\chi})^{\frac{\eta_g-1}{\eta_g}}}. \tag{C.60}$$

Adding these two equations shows that the total tax rate is constant:

$$\tau_{di}^* + \tau_{si}^* = \alpha_g. \tag{C.61}$$

Thus, local population affects the composition of amenity spending, not the total tax burden. Dividing (C.60) by (C.59) gives

$$\frac{\tau_{si}^*}{\tau_{di}^*} = \left(\frac{\nu_g}{1-\nu_g}\right)^{\eta_g} N_i^{-\chi(1-\eta_g)}. \tag{C.62}$$

When $\eta_g < 1$, larger places optimally devote a smaller tax share to shared-capacity amenities because each dollar of such spending is more productive there. However, the service-flow ratio still rises with population:

$$\frac{\tau_{si}^*N_i^{\chi}}{\tau_{di}^*} = \left(\frac{\nu_g}{1-\nu_g}\right)^{\eta_g} N_i^{\chi\eta_g}. \tag{C.63}$$

Large places are therefore relatively abundant in shared-capacity amenities, while small places are relatively scarce in them.

Scale dependence. Define the maximized amenity term

$$M_i(N_i) \equiv (1-\tau_{di}^*-\tau_{si}^*)^{1-\alpha_g} \left((1-\nu_g)\tau_{di}^{*\frac{\eta_g-1}{\eta_g}} + \nu_g(\tau_{si}^*N_i^{\chi})^{\frac{\eta_g-1}{\eta_g}} \right)^{\frac{\alpha_g\eta_g}{\eta_g-1}}. \tag{C.64}$$

Indirect utility in location i is then $A_iN_i^{\sigma}M_i(N_i)$. Population shares solve

$$\frac{N_i}{N} = \frac{(A_iM_i(N_i))^{\frac{1}{\theta-\sigma}}}{\sum_j (A_jM_j(N_j))^{\frac{1}{\theta-\sigma}}}. \tag{C.65}$$

Differentiating this expression with respect to total population gives

$$\frac{d\log N_i}{d\log N} = \frac{1}{\theta - \sigma - \varepsilon_i^g(N_i)}\Gamma, \qquad \varepsilon_i^g(N_i) \equiv \frac{\partial \log M_i(N_i)}{\partial \log N_i}, \tag{C.66}$$

where Γ is a normalizing constant, identical across locations. By the envelope theorem,

$$\varepsilon_i^g(N_i) = \alpha_g \chi \frac{\nu_g(\tau_{si}^* N_i^\chi)^{\frac{\eta_g - 1}{\eta_g}}}{(1-\nu_g)\tau_{di}^{*\,\frac{\eta_g-1}{\eta_g}} + \nu_g(\tau_{si}^* N_i^\chi)^{\frac{\eta_g-1}{\eta_g}}}. \tag{C.67}$$

This expression is the amenity analogue of the housing-supply elasticity in the simple model. When $\eta_g < 1$, the exponent $(\eta_g - 1)/\eta_g$ is negative. Since (C.63) is increasing in N_i, the fraction in (C.67) is strictly decreasing in N_i. Hence $\varepsilon_i^g(N_i)$ is strictly decreasing in local population.

Proposition 7 (Scale dependence from multiple amenities)**.** *Assume $\chi \in (0,1)$, $\eta_g < 1$, and $\theta > \sigma + \alpha_g\chi$. If regions are the same size ($L_i = L$), then a decline in total population N causes an increase in spatial concentration.*

Proof. Since $\varepsilon_i^g(N_i)$ is strictly decreasing in N_i, equation (C.66) implies that $d\log N_i/d\log N$ is strictly decreasing in local population. With equal land sizes, this means the response is strictly decreasing in density N_i/L. The result then follows from Lemma 1. □

The intuition is the same as in the housing and multiple-goods examples. Shared-capacity amenities are abundant in high-density places and scarce in low-density places. When they are complementary to divisible amenities, the marginal value of local scale is highest in small places, where shared-capacity amenities are the weak link. As aggregate population falls, small places lose relatively more of that local scale advantage, so population shares shift toward high-density places.

If instead $\chi = 0$, shared-capacity amenities are no more scalable than divisible amenities, so amenities add no scale force. Or if instead $\eta_g = 1$, the amenity composite is Cobb-Douglas and the scale term becomes a common power of N_i, just like the power amenities in Appendix C.1. In either case, the amenity block affects the level of concentration but does not generate scale dependence.

Appendix D. Model solution algorithms

This appendix describes the numerical algorithms used to solve the dynamic model in Section 4. The key numerical objects are regional population shares, $S_{it} \equiv N_{it}/N_t$, construction employment shares, $\hat{N}_{iXt} \equiv N_{iXt}/N_{it}$, the stock of structures, X_{it}, and expected value functions, $\widehat{V}_{it}$. Given these objects, wages, house prices, and migration probabilities follow from the equilibrium conditions in Section 4.1.

There are several iterations that occur within both of our solution algorithms described below. For some of these iterations, the updates are damped, meaning that we use the following updating equation for the variable on which we are iterating

$$x^{m+1} = \lambda x^{\text{new}} + (1 - \lambda)x^m, \tag{D.1}$$

where x^{new} is the full update and λ is the learning rate in Panel A of Table D.1. We use damping because the fixed-point map for some blocks, especially the transition path construction share block, can have negative derivatives with respect to the previous iterate, which makes undamped iterations oscillate around the fixed point. Table D.1 reports the learning rates and convergence tolerances that we use in our solution algorithms.

Table D.1. Numerical parameters

Algorithm block	Object	Value	Details
Panel A: Learning rates			
Steady state population shares	Population share update	0.20	Damped outer update
Steady state value functions	Value function update	1.00	Full update
Steady state construction shares	Construction share bisection	1.00	Full bisection solution
Transition path population shares	Population share path update	0.50	Damped outer update
Transition path construction shares	Construction share path update	0.10	Damped inner update
Transition path construction bisection	Construction share bisection	1.00	Full bisection solution
Panel B: Convergence tolerances			
Steady state population shares	Population share error	10^{-5}	$\max_i \lvert S_i^{\text{new}} - S_i^m \rvert$
Steady state value functions	Relative expected value function error	10^{-5}	$\max_i \lvert \widehat{V}_i^{\text{new}} - \widehat{V}_i^m \rvert / \max\{\lvert \widehat{V}_i^m \rvert, 10^{-5}\}$
Construction share bisection	Bisection bracket width	10^{-6}	Bracket width for $\hat{N}_{iX}$
Transition path population shares	Population share path error	10^{-5}	$\max_{i,t} \lvert S_{it}^{\text{new}} - S_{it}^m \rvert$
Transition path construction shares	Construction share path error	10^{-5}	$\max_{i,t} \lvert \hat{N}_{iXt}^{\text{new}} - \hat{N}_{iXt}^m \rvert$

Notes: Learning rates are the weights on the full update in (D.1); a value of one means that the block is not damped. Convergence errors are computed using the full update before applying damping.

D.1 Steady state

For a given aggregate population level N, the steady state algorithm iterates on regional population shares. The construction wage can be written as a function of the construction employment share by

inverting the occupation-share equation:

$$w_{iXt}(\hat{N}_{iXt}) = \left(\frac{\hat{N}_{iXt}}{1 - \hat{N}_{iXt}} \frac{z_Y}{z_X} \right)^{\varrho} w_{iYt}. \tag{D.2}$$

Given population shares, the construction block is static: structure accumulation implies

$$X_i = \frac{A_{X,i}}{\delta} \hat{N}_{iX} N_i, \tag{D.3}$$

and substituting (D.3) into the construction Euler equation (44) gives one nonlinear equation in $\hat{N}_{iX}$ for each region. We solve these equations by bisection on $\hat{N}_{iX} \in [0, 1]$, using the midpoint as the first guess for the first region and then using the first region's solution as the guess for the remaining regions. Given construction shares, structures, wages, and house prices, we solve the expected value functions by iterating on (40). We then construct the matrix of migration probabilities, P, and solve directly for the stationary distribution:

$$(I - P')S = 0, \qquad \sum_i S_i = 1, \tag{D.4}$$

replacing one row of the linear system with the normalization. These steps are summarized in Algorithm 1.

Algorithm 1 Steady state solution algorithm

1: Guess an initial population distribution $\{S_j\}$.
2: **repeat**
3: Given $\{S_j\}$ (and hence $\{N_j = S_j N\}$) and $\{\hat{N}_{jX}\}$, compute $\{w_{jY}, w_{jX}, \bar{w}_j, H_j/N_j\}$ using (33), (D.2), and (36), with $N_{jY} = N_j(1 - \hat{N}_{jX})$ and $X_j = (A_{X,j}/\delta)\hat{N}_{jX} N_j$.
4: Solve (44) for $\{\hat{N}_{jX}\}$ using bisection, with tolerance in Table D.1.
5: Solve for $\{\widehat{V}_j\}$ by iterating on (40), using the value function tolerance in Table D.1.
6: Given $\{\widehat{V}_j\}$, compute a new population distribution $\{S_j^{\text{new}}\}$ as the unique solution of (D.4), with transition probabilities from (42).
7: If the population share error is above tolerance, update $\{S_j\}$ using (D.1) with the steady state population share learning rate in Table D.1.
8: **until** $\{S_j\}$ has converged.

D.2 Transition path

We solve transition paths between an initial steady state and a terminal steady state with perfect foresight. We first initialize the path of population shares and construction shares by geometric

interpolation between the initial and terminal steady states. Given a candidate path of construction employment shares, structures evolve forward according to (35). Letting $\xi_{it} \geq 0$ denote the multiplier on the non-negativity constraint $X_{it} \geq (1-\delta)X_{i,t-1}$ in the construction firm's problem, combining the first-order condition with the envelope condition gives

$$\begin{aligned} w_{iXt} = A_{X,i} p_{it} \frac{\partial H(X_{it}, L_i)}{\partial X_{it}} + \beta_X (1-\delta) w_{iX,t+1} \\ - \beta_X (1-\delta) A_{X,i} \xi_{i,t+1} + A_{X,i} \xi_{it}. \end{aligned} \tag{D.5}$$

Defining $q_{it} \equiv w_{iXt} - A_{X,i}\xi_{it}$, this construction Euler equation can be written as

$$q_{it} = A_{X,i} p_{it} \frac{\partial H(X_{it}, L_i)}{\partial X_{it}} + \beta_X (1-\delta) q_{i,t+1}. \tag{D.6}$$

Starting from the terminal steady state value $q_{i,T+1} = w_{iX,T+1}$, we solve (D.6) backward. We then update construction shares by solving $w_{iXt}(\hat{N}_{iXt}) = q_{it}$ at each region and date using (D.2); if $q_{it} \leq 0$, we set construction employment to its lower bound. Given the converged construction block, we solve expected value functions backward from the terminal steady state value, then simulate population shares forward from the initial steady state distribution using (43). These steps are summarized in Algorithm 2.

In the quantitative exercises, one model period is one year and the transition path has $T = 500$ annual periods between the initial and terminal steady states. We set the initial steady state population to $N_0 = 126{,}925{,}843$, Japan's population in 2000, and set the terminal steady state population to $N_T = 10{,}000{,}000$. The aggregate population path is held fixed at N_0 for the first $T_{\text{delay}} = 20$ periods after the path is announced, then declines geometrically with growth factor $(N_T/N_0)^{1/(T_{\text{grow}}+1)}$ and $T_{\text{grow}} = 175$. This path implies that T_{grow} falls dates between the first decline date and the date immediately before the terminal population date; the next date is set to N_T, after which aggregate population is fixed at N_T through the end of the finite horizon. We choose the horizon length to be sufficiently long that the economy reaches the terminal steady state smoothly before the endpoint.

Algorithm 2 Transition path solution algorithm

Solve the initial and terminal steady states using Algorithm 1.
Guess an initial path of population shares $\{S_{jt}\}_{t=1}^{T}$, where T is the length of the transition.
repeat
    Given $\{S_{jt}\}$ (and hence $\{N_{it} = S_{it}N_t\}$), guess an initial path $\{\hat{N}_{iXt}\}$.
    repeat
        Solve forward for $\{X_{it}\}$ using (35) and compute $\{w_{iYt}, w_{iXt}, \bar{w}_{it}, H_{it}/N_{it}, \partial H_{it}/\partial X_{it}\}$ using (33), (D.2), and (36).
        Solve (D.6) backward from $q_{i,T+1} = w_{iX,T+1}$ to obtain $\{q_{it}\}$.
        Update $\hat{N}_{iXt}$: for each i, t, if $q_{it} > 0$, solve $w_{iXt}(\hat{N}_{iXt}) = q_{it}$ using (D.2); if $q_{it} \leq 0$, set $\hat{N}_{iXt} = 0$.
        If the construction share error is above tolerance, damp the update using (D.1) with the transition path construction share learning rate in Table D.1.
    until $\{\hat{N}_{iXt}\}$ has converged.
    Given converged $\{\hat{N}_{iXt}, X_{it}\}$ and all derived quantities, solve (40) backward from the terminal steady state value to obtain $\{\widehat{V}_{jt}\}$.
    Given $\{\widehat{V}_{jt}\}$, solve forward for a new path of population shares $\{S_{jt}^{\text{new}}\}$ using (43).
    If the population share error is above tolerance, damp the update using (D.1) with the transition path population share learning rate in Table D.1.
until $\{S_{jt}\}$ has converged.

D.3 Optimal policy

The optimal policy problem in (45) restricts the planner to the transfer rule in (46) and solves problems at the population level in the initial steady state, N_0, and the terminal steady state, N_T. For numerical implementation, it is convenient to write post-transfer expenditure as

$$x_{io} = (1 + \tau_w)w_{io} + \tau_0, \tag{D.7}$$

where the intercept τ_0 combines the common profit rebate and the uniform component required to make the wage-indexed transfer budget balanced. The steady state equilibrium is computed using Algorithm 1, except that household-side objects use post-transfer expenditures instead of wages. In particular, the occupation shares in (41) replace w_{io} with x_{io}, housing demand is $h_{io} = \alpha x_{io} p_i^{-\zeta_p}$, and the value function and migration probabilities replace $\sum_o z_o w_{jo}^{1/\varrho}$ with $\sum_o z_o x_{jo}^{1/\varrho}$. For a fixed τ_w

and aggregate population N, we solve for τ_0 using the budget residual

$$B(\tau_0; \tau_w, N) \equiv \tau_0 + \tau_w \bar{w}(\tau_0, \tau_w; N) - \bar{\pi}(\tau_0, \tau_w; N) = 0, \tag{D.8}$$

where $\bar{w}$ is the aggregate wage bill per capita and $\bar{\pi}$ is aggregate construction-sector profits per capita in the steady state induced by (τ_w, τ_0). This equation is equivalent to imposing the zero-sum transfer constraint in (45) after adding the common profit rebate. To solve (D.8), we first try fixed-point iteration; if that does not converge, we perform a bisection. Given this budget-balancing intercept, we then solve for the τ_w that maximizes

$$\mathcal{W}(\tau_w; N) = \sum_i S_i(\tau_w; N) \widehat{V}_i(\tau_w; N), \tag{D.9}$$

where the steady state is solved using Algorithm 1. We maximize (D.9) over τ_w using golden-section search. Table D.2 reports the numerical parameters used in this algorithm, and Algorithm 3 summarizes the steps.

Table D.2. Numerical parameters in the optimal policy algorithm

Algorithm block	Numerical parameter	Value
Wage-indexed transfer search	Search interval for τ_w	$[-0.98, 0.98]$
Government budget balance	Initial bracket half-width	$\max\{10^{-4}, 0.02\lvert\tau_0^{(0)}\rvert\}$
Government budget balance	Bracket expansion factor	2
Wage-indexed transfer search	Golden-section tolerance, ϵ_τ	10^{-6}
Government budget balance	Budget residual tolerance, ϵ_B	10^{-6}

Notes: This table reports the numerical parameters used to solve the optimal steady state transfer problem in Algorithm 3.

Algorithm 3 Optimal policy solution algorithm

for $N \in \{N_0, N_T\}$ **do**
  Solve the no-transfer steady state using Algorithm 1.
  Set $\tau_w = 0$ and solve (D.8) for the common profit rebate τ_0; store the resulting no-transfer welfare.
  Define the scalar objective $\mathcal{W}(\tau_w; N)$ in (D.9).
  repeat
    For the candidate value of τ_w proposed by golden-section search, solve (D.8) for τ_0 by fixed-point iteration and, if needed, bisection.
    Given (τ_w, τ_0), solve the induced steady state using Algorithm 1 and compute $\mathcal{W}(\tau_w; N)$.
    Update the golden-section bracket using the welfare values at the candidate points.
  until the bracket for τ_w has converged.
end for

D.4 Software and parallelization

We implement the solution algorithms in `C++17` and compile them with `OpenMP`. The parallel work is over independent region and region-by-date calculations: in the steady state algorithm, we parallelize the value function update, the regional construction-share bisections after the first warm-started region, and construction of the transition-probability system; along transition paths, we parallelize the construction-share updates across regions and dates, as well as the forward population-share updates across regions.

Appendix E. Model estimation

This appendix gives the numerical details behind the estimation procedure summarized in Section 4.2. Let S_i^{data} denote the population share of prefecture i in 2000. We impose $A_{X,i} = 1$ for all i, write $A_i = \bar{A}\tilde{A}_i$ with $\sum_i \tilde{A}_i = 1$, and estimate

$$\Theta = \left(\kappa, \varrho, \delta, \nu, z_Y/z_X, \eta, \alpha, \bar{A}\right), \tag{E.1}$$

where $z_X = 1$ and $\bar{A} = \sum_i A_i$. We set $\{L_i\}$ equal to each prefecture's share of total national inhabitable land, scaled so that $\sum_i L_i = N_0$. Coupled with the normalization $A_{X,i} = 1$, this ensures that structures and land are measured in comparable units. When housing production is Cobb-Douglas, η is fixed at one and removed from Θ.

Our estimation procedure iterates between two steps. First, for a candidate value of Θ, we solve the initial steady state with $N = N_0$ and invert the relative productivities $\{\tilde{A}_i\}$ so that the model matches the prefecture population shares in 2000. Second, holding these inverted fundamentals fixed, we use a Newton method to update Θ so that the model matches the aggregate moments in Table 4. Algorithm 4 summarizes these steps, and the subsections below first describe the inversion of fundamentals and then the construction and matching of target moments. Table E.1 reports the numerical parameters that are used in this algorithm.

Algorithm 4 Estimation algorithm

1: Initialize $\{\tilde{A}_i\}$, set $A_{X,i} = 1$ for all i, choose an initial Θ, and load the target moments.
2: **repeat**
3: Holding Θ fixed, solve the initial steady state at $N = N_0$ using Algorithm 1.
4: Update $\{\tilde{A}_i\}$ using (E.2), reducing Δ_m if the inversion error rises, until the population shares are within the population-share tolerance ϵ_S.
5: Holding the inverted fundamentals fixed, use the trust-region Newton method to choose Θ^{Newton} so that the residuals in (E.18) are within the residual tolerance ϵ_R.
6: Re-solve the initial steady state at Θ^{Newton} and compute both the moment error and the population share inversion error.
7: If both errors are below tolerance, stop. Otherwise, update $\Theta \leftarrow \lambda_\Theta \Theta^{\text{Newton}} + (1 - \lambda_\Theta)\Theta$, where λ_Θ is the Newton update learning rate, and repeat.
8: **until** the moment and inversion errors are both below tolerance.
9: Save the estimated parameter vector and inverted fundamentals.

Table E.1. Numerical parameters in the estimation algorithm

Algorithm block	Numerical parameter	Value
Panel A: Learning rates		
Fundamentals inversion	Initial learning rate, Δ_0	0.50
Fundamentals inversion	Minimum learning rate, $\underline{\Delta}$	0.05
Outer parameter update	Newton update learning rate, λ_Θ	0.15
Panel B: Convergence tolerances		
Fundamentals inversion	Population-share tolerance, ϵ_S	0.002
Moment matching	Residual tolerance, ϵ_R	0.01
Moment matching	Residual scaling floor, ϵ_0	10^{-15}
Newton step	Scaled step tolerance, ϵ_d	10^{-8}
Newton step	Scaled gradient tolerance, ϵ_g	0.01
Panel C: Trust-region controls		
Newton step	Finite-difference relative step	1.49×10^{-8}
Newton step	Finite-difference step floors	0.1 or 0.001
Newton step	Acceptance threshold	0.15
Newton step	Initial trust-region radius	1.00
Newton step	Maximum trust-region radius	1000

Notes: This table reports the values of numerical parameters that are used in Algorithm 4.

E.1 Inversion of fundamentals

For each candidate value of Θ, we invert the relative goods productivities $\{\tilde{A}_i\}$ so that the model's steady state matches prefecture-level population shares in 2000. After solving for the steady state using Algorithm 1, we update the relative productivity terms using

$$\tilde{A}_i^{m+1} = \tilde{A}_i^m \left(\frac{S_i^{\text{data}}}{S_i^{\text{model}}} \right)^{\Delta_m}, \qquad \sum_i \tilde{A}_i^{m+1} = 1. \tag{E.2}$$

We initialize the learning rate at $\Delta_m = \Delta_0$, where Δ_0 is the initial learning rate. If the population-share inversion error rises from one iteration to the next, we halve the learning rate down to the floor $\underline{\Delta}$, the minimum learning rate, and we stop the inversion when the maximum absolute population-share error is below the population-share tolerance ϵ_S.

E.2 Target moments

We compute each model moment listed in Table 4 in the initial steady state with $N = N_0$, after applying the productivity inversion described in the previous subsection.

Average out-migration rate. The first moment is the average out-migration rate, which we compute from the model-implied migration probabilities as

$$\sum_i S_i(1 - \pi_{ii}), \tag{E.3}$$

where π_{ii} is the probability that a household in prefecture i remains in prefecture i between periods. We compute the empirical counterpart as the population-weighted average across prefectures of annual out-migrants among Japanese residents divided by the Japanese population, averaged over 1990–2000 using the e-Stat Statistics Dashboard prefecture data.

Housing supply elasticity. The second moment is the housing supply elasticity, which we derive in the model as follows. The steady state landlord condition implies an inverse supply curve for housing,

$$A_{X,i} p_i \frac{\partial H_i}{\partial X_i} = \left[1 - \beta_X (1 - \delta)\right] w_{iX}. \tag{E.4}$$

Log-differentiating both sides with respect to X_i, holding N_i and L_i fixed, gives

$$\frac{d \log p_i}{d \log X_i} + \frac{d \log(\partial H_i / \partial X_i)}{d \log X_i} = \frac{d \log w_{iX}}{d \log X_i}. \tag{E.5}$$

The structures elasticity in housing production is

$$s_{Xi} \equiv \frac{\partial H_i}{\partial X_i} \frac{X_i}{H_i} = \frac{(1-\nu) X_i^{1-1/\eta}}{(1-\nu) X_i^{1-1/\eta} + \nu L_i^{1-1/\eta}}, \tag{E.6}$$

so $\partial H_i / \partial X_i = (1-\nu)(H_i / X_i)^{1/\eta}$ implies

$$\begin{aligned} \frac{d \log(\partial H_i / \partial X_i)}{d \log X_i} &= \frac{1}{\eta} \left(\frac{d \log H_i}{d \log X_i} - 1 \right) \\ &= \frac{s_{Xi} - 1}{\eta}. \end{aligned} \tag{E.7}$$

For the right-hand side of (E.5), steady state structures satisfy $X_i = (A_{X,i}/\delta) \hat{N}_{iX} N_i$. Holding N_i fixed, the construction wage can therefore be written as

$$\begin{aligned} \log w_{iX} = \varrho \log \left(\frac{\hat{N}_{iX}}{1 - \hat{N}_{iX}} \right) \\ + \sigma \log \left((1 - \hat{N}_{iX}) N_i \right) + \text{constant}, \end{aligned} \tag{E.8}$$

which implies

$$\frac{d \log w_{iX}}{d \log X_i} = \frac{\varrho - \sigma \hat{N}_{iX}}{1 - \hat{N}_{iX}}. \tag{E.9}$$

Substituting these terms into (E.5) gives

$$\frac{d \log p_i}{d \log X_i} = \frac{\varrho - \sigma \hat{N}_{iX}}{1 - \hat{N}_{iX}} + \frac{1 - s_{Xi}}{\eta}. \tag{E.10}$$

Finally, converting from structures to housing using $d \log H_i = s_{Xi} d \log X_i$ gives the housing supply elasticity

$$\frac{d \log H_i}{d \log p_i} = \frac{\eta s_{Xi}(1 - \hat{N}_{iX})}{\eta(\varrho - \sigma \hat{N}_{iX}) + (1 - s_{Xi})(1 - \hat{N}_{iX})}. \tag{E.11}$$

The numerator reflects the response of the housing stock to additional structures, while the denominator combines the labor reallocation margin, governed by ϱ, and diminishing returns from fixed land, governed by $1 - s_{Xi}$. We evaluate (E.11) in Tokyo to construct the model moment in the second row of Table 4.

Housing depreciation rate. With land fixed and structures depreciating at rate δ, region i's housing stock evolves according to

$$\frac{dH_i}{dt} = \frac{\partial H_i}{\partial X_i} \frac{dX_i}{dt} = (1 - \nu) \left(\frac{H_i}{X_i} \right)^{1/\eta} (-\delta X_i) = -s_{Xi} \delta H_i. \tag{E.12}$$

The local housing depreciation rate is therefore $\delta_{h,i} = s_{Xi} \delta$. We aggregate these local rates across regions using housing-stock weights:

$$\frac{\sum_i H_i s_{Xi} \delta}{\sum_i H_i}. \tag{E.13}$$

In the third row of Table 4, we match this aggregate rate to Japan's aggregate housing depreciation rate, which we compute from the Economic and Social Research Institute (ESRI) national accounts as gross fixed capital formation in dwellings minus the change in real net housing capital, divided by the beginning-of-year real net housing capital stock and averaged over 1995–2024.

Construction factor income and employment-to-income ratio. Our fourth moment is the construction factor income share. In the model, construction labor is the only non-land factor in the construction sector, so the model counterpart is

$$\frac{\sum_i S_i \hat{N}_{iX} w_{iX}}{\sum_i S_i \left[p_i h_i + (1 - \hat{N}_{iX}) w_{iY} \right]}. \tag{E.14}$$

This moment is primarily informative about ν because ν changes the share of housing value paid to fixed land rather than construction labor. We compute the data counterpart from ESRI national accounts as construction domestic factor income divided by aggregate domestic factor income in 2000, where domestic factor income equals gross output minus intermediate inputs, consumption of fixed capital, and taxes on production and imports less subsidies. The fifth moment is the

construction employment share divided by the construction factor income share,

$$\frac{\sum_i S_i \hat{N}_{iX}}{\left[\sum_i S_i \hat{N}_{iX} w_{iX}\right] \Big/ \left[\sum_i S_i \left[p_i h_i + (1 - \hat{N}_{iX}) w_{iY}\right]\right]}. \tag{E.15}$$

Since the construction factor income share is matched in the previous row, this ratio is equivalent to matching the aggregate construction employment share while using the scale-free statistic reported in Table 4. However, the ratio is more directly informative about $\frac{z_Y}{z_X}$, which is why we choose to target it. We compute the construction employment share in the data as construction employed persons divided by total employed persons in the ESRI national accounts in 2000, and then divide it by the construction factor income share to form the target ratio.

Housing supply elasticity vs. density slope. The sixth moment uses the same regional elasticities in (E.11) to target the slope of housing supply elasticity with respect to population density. We construct this moment as the unweighted OLS slope from regressing ε_i^H on $\log(N_i/L_i)$ across prefectures in the initial steady state, which disciplines η because η governs how the land share, and hence the local supply elasticity, varies with density. The data counterpart is the slope reported in the left panel of Figure 4.

Housing expenditure share. The seventh moment is the housing expenditure share

$$\frac{\sum_i S_i p_i h_i}{\sum_i S_i \left[\hat{N}_{iX} w_{iX} + (1 - \hat{N}_{iX}) w_{iY}\right]}. \tag{E.16}$$

We match this moment to a standard value of 0.25 from Piazzesi and Schneider (2016).

Aggregate GDP per capita. The final moment is aggregate GDP per capita, which we normalize to one and compute as

$$\sum_i S_i \left[p_i h_i + (1 - \hat{N}_{iX}) w_{iY}\right]. \tag{E.17}$$

The first term is housing-sector value added per capita, and the second term is goods-sector value added per capita. This normalization pins down the aggregate productivity scale $\bar{A} = \sum_i A_i$ after the relative productivities $\{\tilde{A}_i\}$ have been inverted to match the spatial distribution of population.

E.3 Matching model and data moments

Holding the inverted fundamentals fixed, we choose Θ to match the target moments from the previous subsection using a trust-region Newton method. The method repeatedly linearizes the mapping from parameters to moment residuals, proposes a parameter update that reduces

the squared residuals in this local approximation, and then limits the step size when the local approximation is poor. The residual for each moment ℓ is

$$R_\ell(\Theta) = \frac{m_\ell^{\text{model}}(\Theta) - m_\ell^{\text{data}}}{\max\{|m_\ell^{\text{data}}|, \epsilon_0\}}. \tag{E.18}$$

where ϵ_0 is the residual scaling floor. We approximate the Newton Jacobian using forward finite differences, with perturbations governed by the finite-difference relative step and finite-difference step floors in Table E.1. We accept trial steps using the acceptance threshold and cap the trust-region radius at the maximum trust-region radius after initializing it at the initial trust-region radius. We stop the Newton subproblem when the maximum absolute residual is below the residual tolerance ϵ_R, when the scaled step norm is below the scaled step tolerance ϵ_d, or when the scaled gradient norm of the squared-residual objective is below the scaled gradient tolerance ϵ_g. If either the moment residuals or the population-share inversion error remain above tolerance after the Newton step, we damp the outer update by setting $\Theta \leftarrow \lambda_\Theta \Theta^{\text{Newton}} + (1 - \lambda_\Theta)\Theta$ and repeat the inversion and moment-matching steps.

Appendix F. Additional results from the dynamic model

F.1 Results discussed in Sections 5.1 and 5.2

Figure F.1. Land share of housing value: data and model

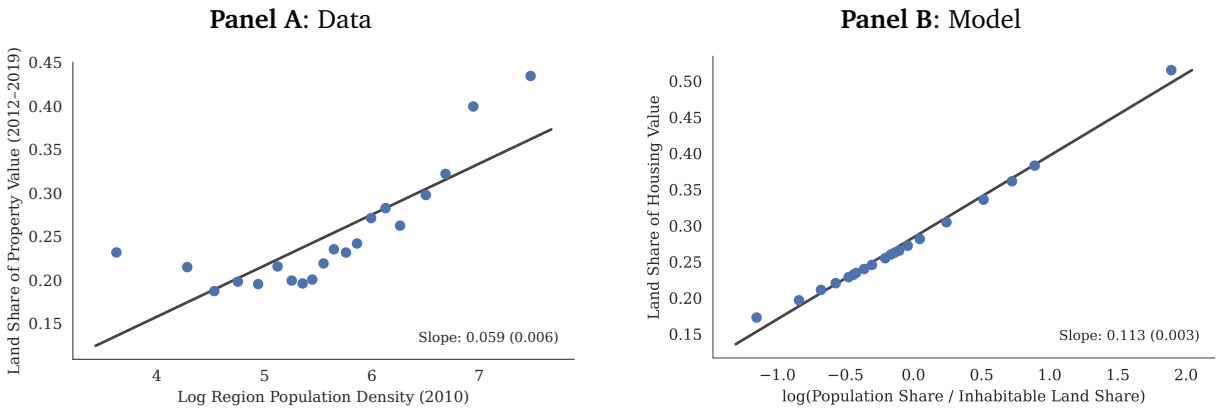

Notes: Panel A reproduces Figure A.8. Panel B plots the corresponding relationship across prefectures in the baseline model's initial steady state. The horizontal axis is the log of model population share divided by inhabitable land share. The model-implied land share is the output elasticity of land in housing production, $1 - s_{Xi}$, where s_{Xi} is defined in (E.6). Both panels use equally-spaced bins of the horizontal variable; fitted lines and reported slopes are estimated on the unbinned observations using OLS with heteroskedasticity-robust standard errors in parentheses. Go back.

Figure F.2. Inverted productivities $\{A_i\}$ across prefectures

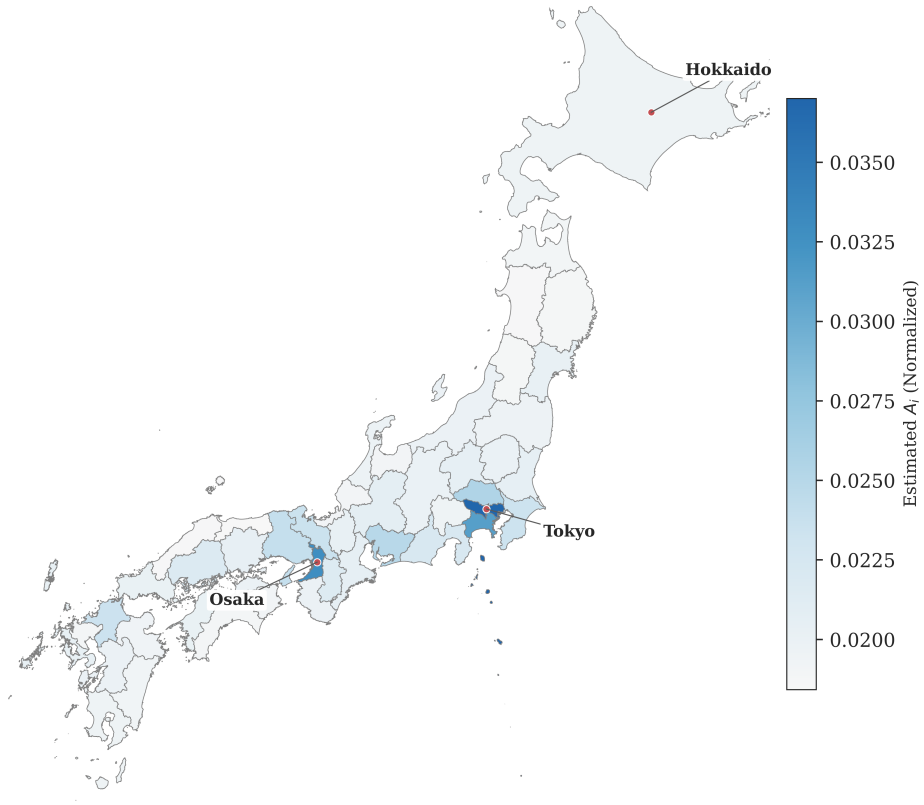

Notes: The figure plots the inverted productivities recovered by matching 2000 prefecture population shares, as described in Appendix E. In this plot, the value of the productivity in each region is normalized by $\sum_i A_i$ so that they sum to one. Go back.

Figure F.3. Model fit on population shares by prefecture

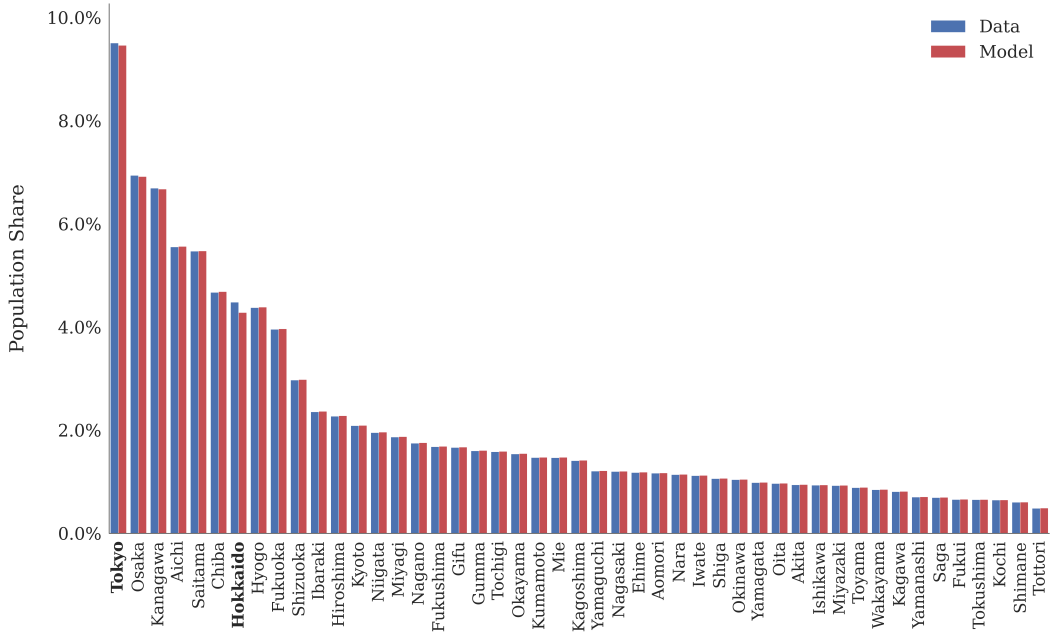

Notes: The figure compares 2000 prefecture population shares with the population shares in the estimated model. Population shares are targeted in the fundamentals inversion described in Appendix E. Go back.

Figure F.4. Inhabitable land area share by prefecture

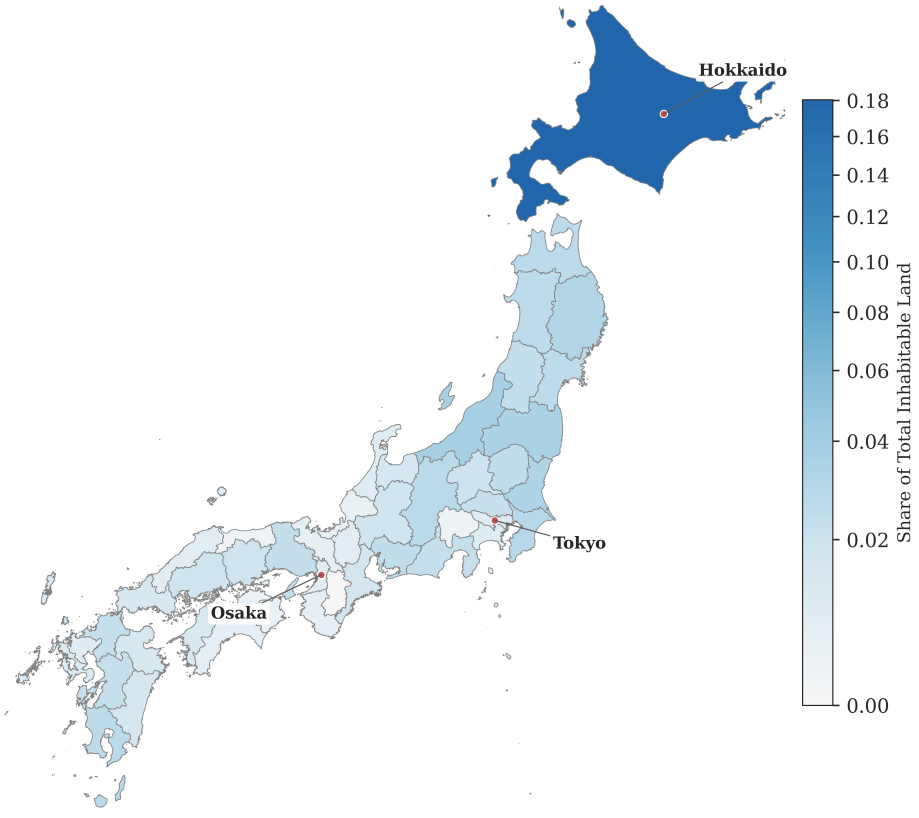

Notes: The figure plots each prefecture's share of national inhabitable land area from the Statistics Bureau of Japan e-Stat Statistics Dashboard. These shares determine the relative land endowments in the model. Go back.

Table F.1. Targeted moments and model fit: Cobb-Douglas housing supply and demand

Parameter	Value	Moment	Source	Model	Data
κ	$10^{6.46}$	Average out-migration rate	Japan 1990-2000	0.024	0.024
ϱ	2.00	Tokyo housing supply elasticity	Baum-Snow and Han (2024)	0.211	0.210
δ	0.10	Housing depreciation rate	Japan 1995-2024	0.058	0.058
ν	0.44	Construction factor income share	Japan 2000	0.085	0.085
z_Y/z_X	9.22	Construction empl. / income	Japan 2000	1.151	1.151
α	0.25	Housing expenditure share	Piazzesi and Schneider (2016)	0.250	0.250
$\sum_i A_i$	17.24	Aggregate GDP per capita	Normalization	1.009	1.000

Notes: This table reproduces the results in Table 4 when estimating an alternative model with Cobb-Douglas housing supply and Cobb-Douglas housing demand, so $\eta = \zeta_p = 1$. The elasticity-density-slope target is omitted because η is fixed. Go back.

Table F.2. Targeted moments and model fit: CES housing supply and Cobb-Douglas housing demand

Parameter	Value	Moment	Source	Model	Data
κ	$10^{6.46}$	Average out-migration rate	Japan 1990-2000	0.024	0.024
ϱ	1.05	Tokyo housing supply elasticity	Baum-Snow and Han (2024)	0.210	0.210
δ	0.09	Housing depreciation rate	Japan 1995-2024	0.058	0.058
ν	0.31	Construction factor income share	Japan 2000	0.085	0.085
z_Y/z_X	9.58	Construction empl. / income	Japan 2000	1.151	1.151
η	0.65	Housing supply vs. density slope	Baum-Snow and Han (2024)	-0.097	-0.097
α	0.25	Housing expenditure share	Piazzesi and Schneider (2016)	0.250	0.250
$\sum_i A_i$	16.86	Aggregate GDP per capita	Normalization	1.002	1.000

Notes: This table reproduces the results in Table 4 when estimating an alternative model with Cobb-Douglas housing demand, $\zeta_p = 1$, while estimating the CES housing-supply parameter η. Go back.

Table F.3. Targeted moments and model fit: Cobb-Douglas housing supply with PIGL housing demand

Parameter	Value	Moment	Source	Model	Data
κ	$10^{6.45}$	Average out-migration rate	Japan 1990-2000	0.024	0.024
ϱ	1.98	Tokyo housing supply elasticity	Baum-Snow and Han (2024)	0.210	0.210
δ	0.10	Housing depreciation rate	Japan 1995-2024	0.058	0.058
ν	0.44	Construction factor income share	Japan 2000	0.085	0.085
z_Y/z_X	8.87	Construction empl. / income	Japan 2000	1.151	1.151
α	0.33	Housing expenditure share	Piazzesi and Schneider (2016)	0.250	0.250
$\sum_i A_i$	16.69	Aggregate GDP per capita	Normalization	1.001	1.000

Notes: This table reproduces the results in Table 4 when estimating an alternative model with Cobb-Douglas housing supply, $\eta = 1$, while retaining PIGL housing demand with $\zeta_p = 0.75$. The elasticity-density-slope target is omitted because η is fixed. Go back.

Figure F.5. Decomposition of the steady state change in GDP per capita

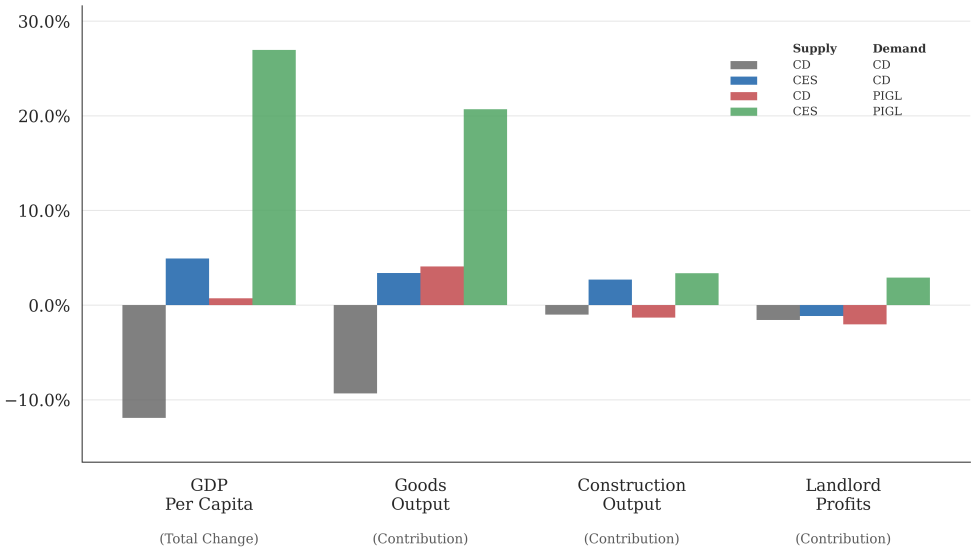

Notes: The figure decomposes the percent change in aggregate GDP per capita between the initial and terminal steady states into the contributions of goods output, construction output, and landlord profits, defined respectively as $\sum_i (N_i/N) w_{iY}(1 - \hat{N}_{iX})$, $\sum_i (N_i/N) w_{iX} \hat{N}_{iX}$, and $\sum_i (N_i/N)\left(p_i h_i - \hat{N}_{iX} w_{iX}\right)$. Each contribution is divided by initial steady state aggregate GDP per capita. Bars correspond to the baseline model (CES housing supply with PIGL housing demand) and the three alternative estimated models described in Section 4.2. Go back.

Figure F.6. Additional changes between steady states

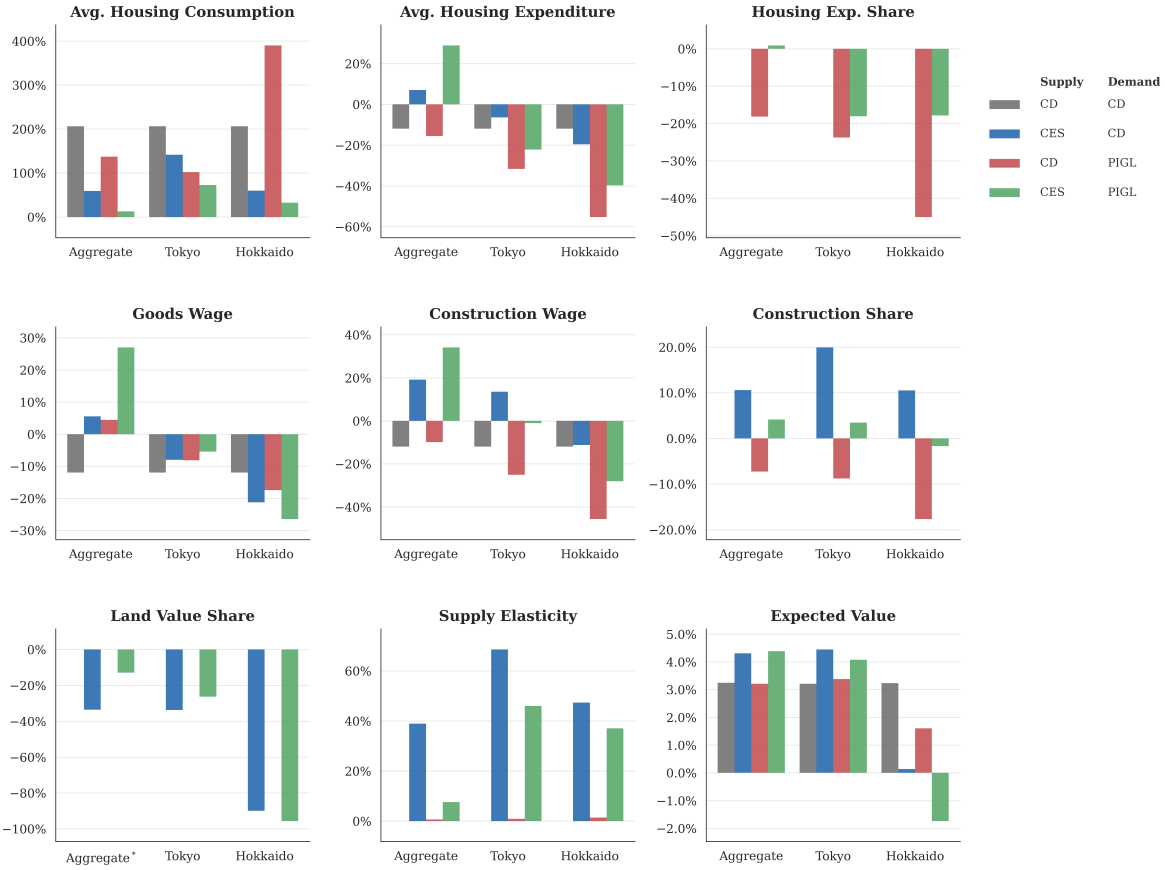

Notes: The figure reports the percent change in nine variables between the initial and terminal steady states, for the aggregate economy, Tokyo, and Hokkaido. Bars correspond to the baseline model (CES housing supply with PIGL housing demand) and the three alternative estimated models described in Section 4.2. The aggregate housing expenditure share is computed as aggregate housing expenditure divided by aggregate post-tax wage income, and the aggregate land value share is the housing-value-weighted mean of the prefecture land share $1 - s_{Xi}$ defined in (E.6). Go back.

Figure F.7. Additional variables at the start of the transition

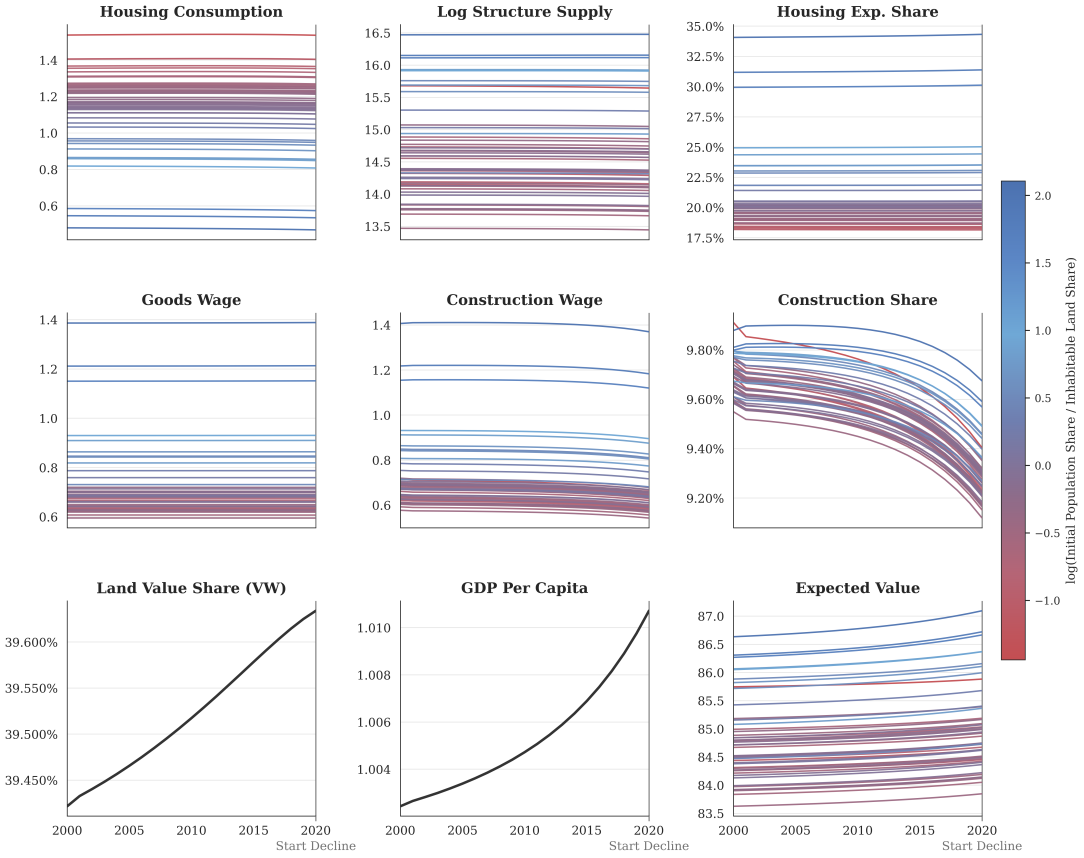

Notes: The figure reports nine variables along the transition path of the baseline model over 2000–2020, when the aggregate population is constant. Each panel plots a separate variable in levels, except structure supply, which is plotted in logs; lines correspond to individual prefectures, color-coded by initial log population share over inhabitable land share, $\log(S_i/L_i)$. The land value share and GDP per capita are aggregate series rather than prefecture-level, with the aggregate land value share computed as the housing-value-weighted mean of the prefecture land share $1 - s_{Xi}$ defined in (E.6). The shaded interval marks the start of the population decline announced in 2000. Go back.

Figure F.8. Additional variables at the end of the transition

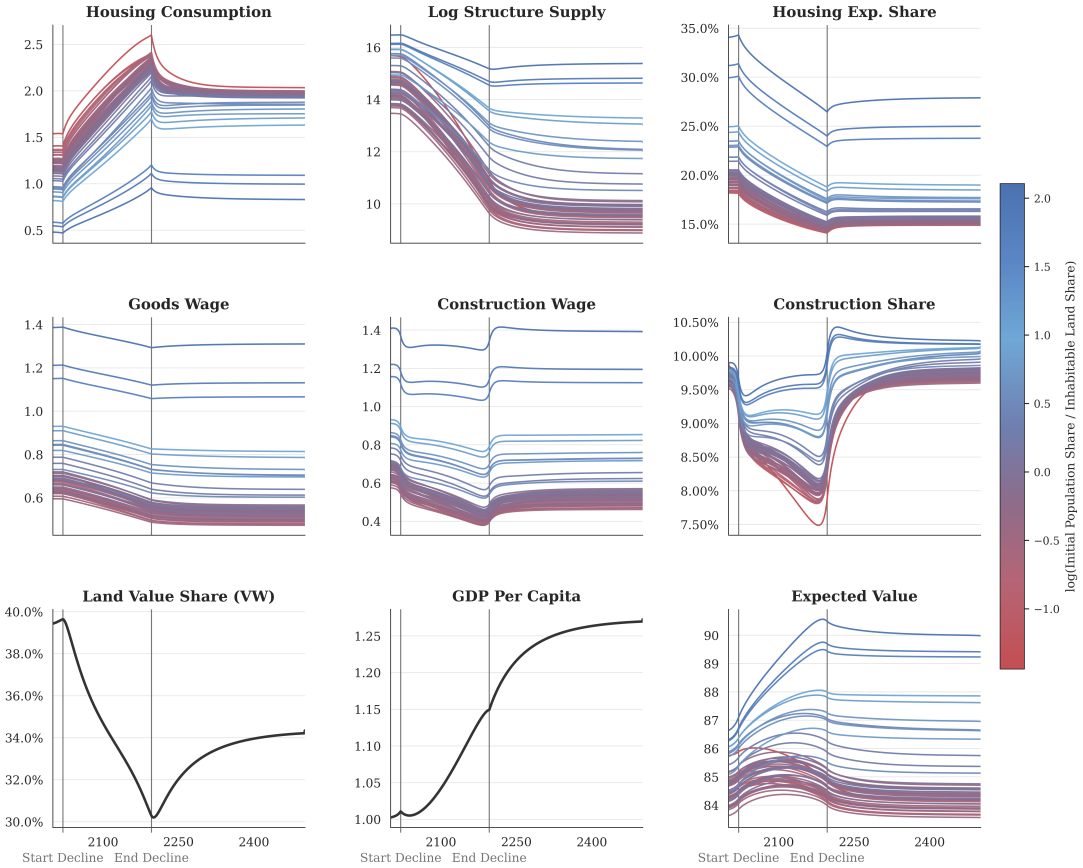

Notes: The figure reports the same level variables as Figure F.7 over the full transition path of the baseline model to the terminal steady state. Vertical markers indicate the start and end of the aggregate population decline. Go back.

Figure F.9. Cumulative net migration to Tokyo by origin prefecture, 2010–2019: data and model

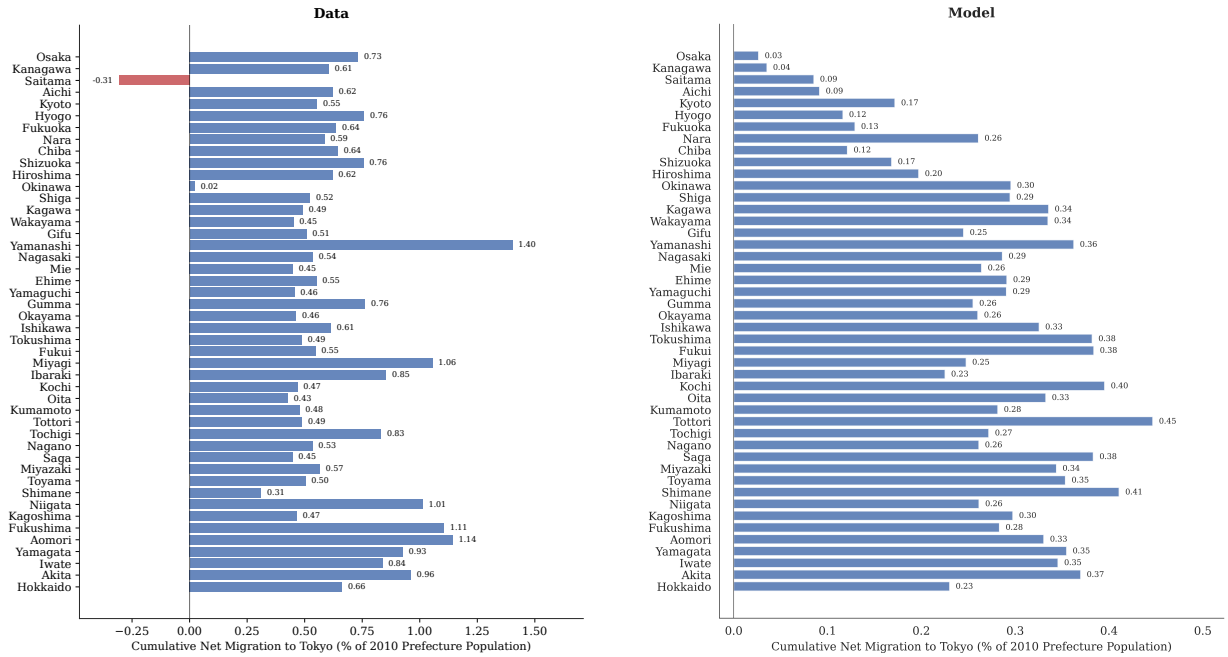

Notes: Each bar reports cumulative net migration to Tokyo from the corresponding non-Tokyo origin prefecture between 2010 and 2019, expressed as a percent of the origin prefecture's 2010 population. The left panel uses Japanese, both-sex prefecture origin-destination flows from the e-Stat Report on Internal Migration in Japan; net migration equals migrants from the origin to Tokyo minus migrants from Tokyo to the origin. The right panel uses the model-implied bilateral migration probabilities $\mathbb{P}(j_{kt} = j \mid j_{k,t-1} = i)$ in (42) along the baseline transition path, multiplied by model origin populations and cumulated over the same years. Prefectures are sorted by 2000 population share over fixed-2010 inhabitable land share. Go back.

F.2 Results discussed in Section 6.2

Figure F.10. Welfare gains from optimal steady state transfers

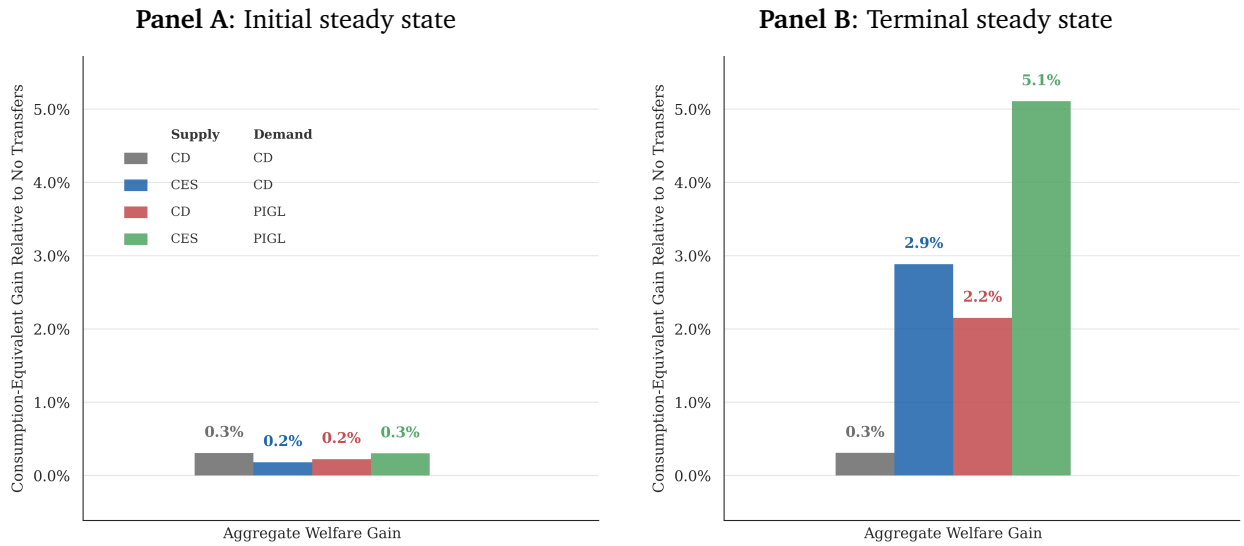

Notes: The figure reports the percent certainty-equivalent welfare gain from the optimal transfer policy relative to the no-transfer baseline in the initial steady state (Panel A) and terminal steady state (Panel B). Welfare gains are computed as $\exp((1-\beta)(W^{\text{SP}} - W^0)) - 1$, where W^{SP} is aggregate welfare under the optimal policy and W^0 is aggregate welfare with no transfers (but profits are still redistributed). Colors index the four estimated models indicated in the legends. Go back.

Figure F.11. Expected value functions across steady states

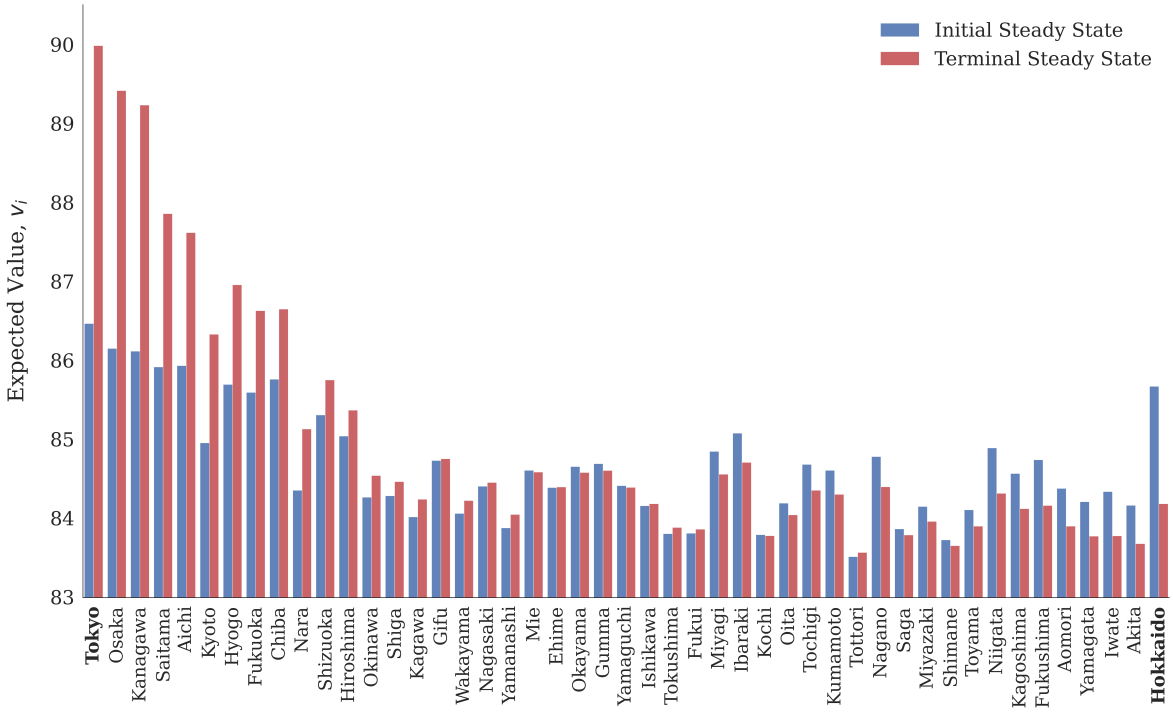

Notes: The figure reports each prefecture's expected value function, $\widehat{V}_i$, from the baseline model without transfers in the initial steady state (blue) and the terminal steady state (red). Prefectures are sorted by their initial population share over inhabitable land share. Go back.

F.3 Extension: spending commitments

Model description. Relative to the baseline model in Table 3, the first change is that each local government taxes labor income at a proportional rate τ_{jt} and returns the proceeds as a local public good. The household budget constraint becomes

$$(1-\tau_{jt})w_{jot} \geq c_{jokt} + p_{jt}h_{jokt}. \tag{F.1}$$

With homothetic PIGL preferences, housing demand is $h_{jot} = \alpha(1-\tau_{jt})w_{jot}p_{jt}^{-\zeta_p}$. Aggregating across occupations and imposing housing-market clearing gives

$$p_{jt} = \left(\alpha \frac{(1-\tau_{jt})\bar{w}_{jt}}{H(X_{jt}, L_j)/N_{jt}}\right)^{1/\zeta_p}, \tag{F.2}$$

where $\bar{w}_{jt} \equiv \sum_o (N_{jot}/N_{jt})w_{jot}$ is the average wage in location j. The occupation shares in (41) are unchanged because taxes, house prices, and public spending are common across occupations within a location.

The second change is households' flow utility, which now depends on per capita spending. We assume that this term enters household utility multiplicatively with elasticity α_g. Therefore, combining (F.1) and (F.2) with the nested Fréchet shocks gives the following expression for the expected value function:

$$\begin{aligned}\widehat{V}_{i,t-1} = \theta \log\Bigg(\sum_j \Bigg(\sum_o z_o w_{jot}^{\frac{1}{\varrho}}\Bigg)^{\frac{\varrho}{\theta}} \left(\frac{1-\tau_{jt}}{\kappa_{ij}}\right)^{\frac{1}{\theta}} \\ \times \exp\left\{-\frac{\alpha}{\theta}\frac{p_{jt}^{1-\zeta_p}-1}{1-\zeta_p}\right\}(\tau_{jt}\bar{w}_{jt})^{\frac{\alpha_g}{\theta}} \exp\left\{\frac{\beta}{\theta}\widehat{V}_{jt}\right\}\Bigg) + \theta\bar{\gamma}.\end{aligned} \tag{F.3}$$

This expression replaces (40). The migration probabilities in (42) and the population-share law of motion in (43) are unchanged in form, but they are computed using the location-specific flow term in (F.3).

The final change is the spending-commitment block. Local governments target total spending equal to a fixed tax share τ^* of the local wage bill,

$$G_{it}^* = \tau^* \bar{w}_{it} N_{it}. \tag{F.4}$$

The spending commitment is a weighted average of this target and the previous period's spending,

$$\tilde{G}_{it} = (1-\phi)G_{it}^* + \phi\tau_{i,t-1}\bar{w}_{i,t-1}N_{i,t-1}. \tag{F.5}$$

The implemented tax rate is the tax rate that finances this spending commitment:

$$\tau_{it} = \frac{\tilde{G}_{it}}{\bar{w}_{it}N_{it}} = (1-\phi)\tau^* + \phi\tau_{i,t-1}\frac{\bar{w}_{i,t-1}N_{i,t-1}}{\bar{w}_{it}N_{it}}. \tag{F.6}$$

Thus, when a prefecture loses population or wage income, its tax rate rises to finance previously committed spending.

Calibration and estimation. We calibrate the additional parameters in this extension to the values listed in Table F.4. First, we set the target steady state tax rate to 10%, matching the standard local inhabitant tax rate in Japan.[11] Second, we set the weight of public spending in utility equal to 0.0741. We choose this value because a planner that chooses a location-specific tax rate in steady state to maximize the weighted sum of flow utilities across the two occupations would choose a tax rate of $\tau^{sp} = \alpha_g/(\alpha_g + 1 - s_{h,i}/\zeta_p)$, where $s_{h,i}$ is the housing expenditure share in region i. Setting $\tau^{sp} = 0.1$ and $s_{h,i}$ equal to the aggregate housing expenditure share for all regions gives $\alpha_g = 0.0741$. Finally, we set the persistence of government spending to 0.9. We choose this value, which is higher than estimates in Japan from Morita (2020), for two reasons. First, we are interested in how the components of government spending that finance local public goods adjust, which are likely more persistent. Second, our interest is in assessing the quantitative importance of this mechanism, and choosing a high ϕ allows us to place an upper bound on its importance. After calibrating these parameters, we then estimate the remaining model parameters following the procedure in Appendix E. Table F.5 reports the estimated parameters and model fit.

Table F.4. Additional parameter values: spending commitments

Parameter		Value
Target steady state tax rate	τ^*	0.10
Spending persistence	ϕ	0.90
Public spending utility weight	α_g	0.0741

Notes: The table reports the additional parameter values used in the spending commitment extension.

Results. This spending-commitment extension introduces two distinct changes to the baseline model: (i) adding local public goods that are financed with proportional taxes; and (ii) making the spending on these goods adjust slowly. Because we are interested in the effects of the latter but the former also affects our results, we highlight the effects of spending commitments by comparing the same re-estimated model with $\phi = 0$ to the model in which ϕ is set to the value in Table F.4.[12] Note

[11] The 10% local inhabitant tax combines prefectural and municipal taxes; see, for example, the Japanese Local Tax Act summary.

[12] The introduction of local public goods has minimal effects on our results, increasing Tokyo's population share in the terminal steady state from 39% to 43%.

Table F.5. Parameter estimates and model fit: spending commitments

Parameter	Value	Moment	Source	Model	Data
κ	$10^{6.45}$	Average out-migration rate	Japan 1990-2000	0.024	0.024
ϱ	1.39	Tokyo housing supply elasticity	Baum-Snow and Han (2024)	0.211	0.210
δ	0.08	Housing depreciation rate	Japan 1995-2024	0.059	0.058
ν	0.17	Construction factor income share	Japan 2000	0.085	0.085
z_Y/z_X	9.10	Construction empl. / income	Japan 2000	1.151	1.151
η	0.60	Housing supply vs. density slope	Baum-Snow and Han (2024)	-0.097	-0.097
α	0.37	Housing expenditure share	Piazzesi and Schneider (2016)	0.250	0.250
$\sum_i A_i$	17.73	Aggregate GDP per capita	Normalization	1.010	1.000

Notes: This table reproduces the results in Table 4 when estimating the model with the spending commitment extension.

that changing ϕ does not require re-estimating the model, since the estimation is done in a steady state and ϕ affects only transition dynamics.

Figure F.12 plots how tax rates evolve in the model with $\phi > 0$ in anticipation of the population decline and along the full transition path. In anticipation of the transition, Panel A shows that less dense prefectures see increases in tax rates, while denser prefectures see decreases. These differences are driven by the migration responses in anticipation of the decline discussed in Section 5.2: (F.6) implies that tax rates move inversely to population growth. Once the population decline begins, Panel B shows that tax rates begin increasing in all prefectures because population levels start declining everywhere. By the end of the population decline, the differences in tax rates are substantial: Hokkaido's tax rate rises from 10% to 15.5%, while Tokyo's remains below 11% for the entire transition.

Figure F.12. Tax rates during the transition: spending commitments

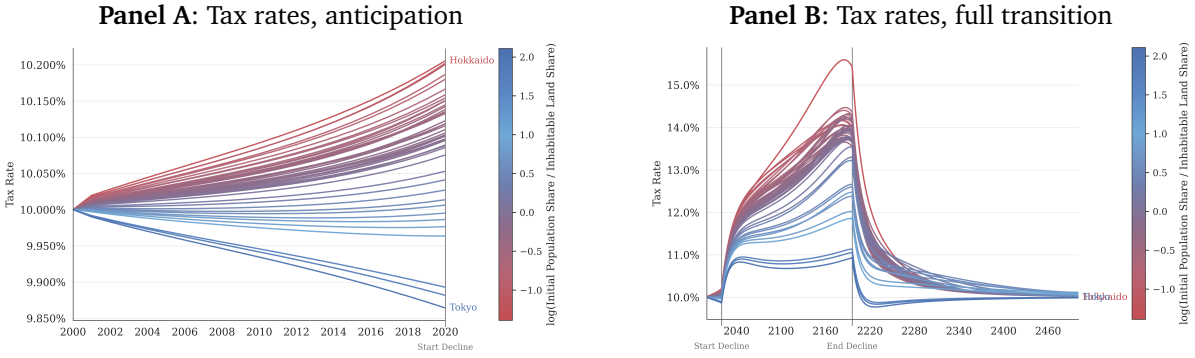

Notes: This figure plots tax rates, computed using (F.6), along the transition in the model with the spending commitment extension. See the notes to Figure 8 for additional details on how this figure is constructed.

While spending commitments have no effect on changes across steady states because tax rates eventually revert back to their original level, it speeds up the migration response in anticipation of the decline. Panel A of Figure F.13 reproduces the anticipation scatter for the public-goods

specification with $\phi = 0$, while Panel B sets ϕ to the value in Table F.4. Comparing the two shows that spending commitments increase the unweighted slope by 16%. This is meaningful, but still small relative to our baseline effect.

Figure F.13. Initial population density and subsequent population growth: spending commitments

Panel A: No spending commitments, $\phi = 0$

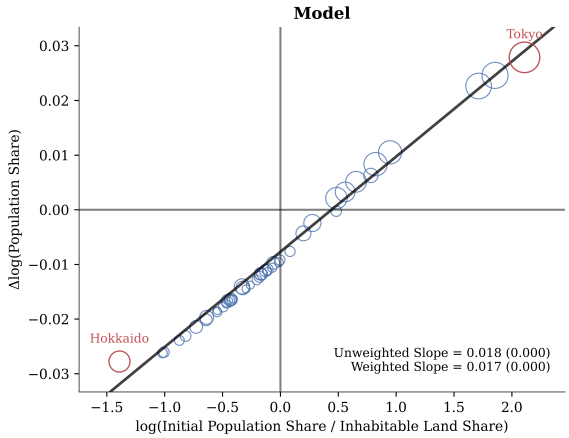

Panel B: Spending commitments, $\phi > 0$

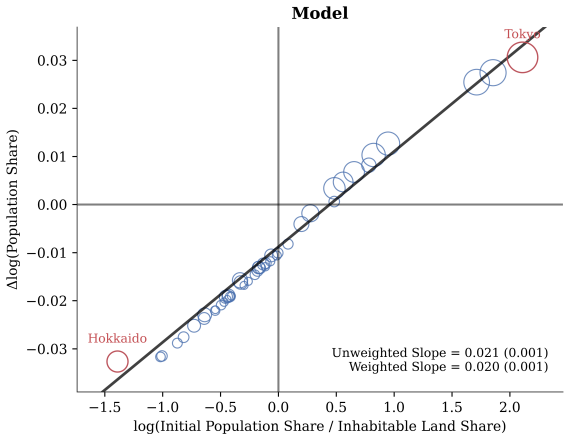

Notes: This figure reproduces Figure 9 in the model with the spending commitment extension with $\phi = 0$ (Panel A) and ϕ set to the value in Table F.4 (Panel B).

F.4 Extension: amenity spillovers

Model description. We incorporate amenity spillovers in the model in Table 3 by introducing a reduced-form term in flow utility that has a constant elasticity with respect to local population, N_{jt}^{γ}, as in Ahlfeldt et al. (2015). Households' value function becomes

$$V_{ikt} = \max_{j,o,c_{jokt},h_{jokt}} \left\{ \log\left(U(c_{jokt}, h_{jokt}) N_{jt}^{\gamma} \epsilon_{jokt} / \kappa_{ij} \right) + \beta \mathbb{E}_t V_{jk,t+1} \right\}. \tag{F.7}$$

Combining (F.7) with the nested Fréchet shocks gives the expected value function

$$\begin{aligned} \widehat{V}_{i,t-1} = \theta \log\Bigg(\sum_j \Bigg(\sum_o z_o w_{jot}^{\frac{1}{\varrho}} \Bigg)^{\frac{\varrho}{\theta}} \left(\frac{N_{jt}^{\gamma}}{\kappa_{ij}} \right)^{\frac{1}{\theta}} \\ \times \exp\left\{ -\frac{\alpha}{\theta} \frac{p_{jt}^{1-\zeta_p} - 1}{1-\zeta_p} \right\} \exp\left\{ \frac{\beta}{\theta} \widehat{V}_{jt} \right\} \Bigg) + \theta \bar{\gamma}. \end{aligned} \tag{F.8}$$

This expression replaces (40). The migration probabilities in (42) and the population-share law of motion in (43) are unchanged in form, but they are computed using the location-specific flow term in (F.8).

Calibration and estimation. The key question is how to calibrate γ, about which there is a wide range of estimates. Fajgelbaum and Gaubert (2020) use the estimates from Diamond (2016) to back out the implied spillovers across workers of different skills. Depending on the relative share of high- vs. low-skilled workers, these estimates imply a negative spillover elasticity between -0.25 and -0.47.[13] In contrast, Ahlfeldt et al. (2015) estimate a positive amenity spillover of $\gamma \approx 0.08$ that declines rapidly with distance, which Giannone et al. (2026) argue is consistent with declines in amenities in depopulating Japanese prefectures, and Ahlfeldt and Pietrostefani (2019) conduct a review of amenity spillover estimates, many of which are positive. Given this range of estimates, we choose a value of $\gamma = -0.18$, which is in the middle of this range. We view this choice as conservative because these estimates are typically at more granular levels of aggregation than a prefecture and (as described below) the degree of scale dependence is increasing in γ. We then re-estimate the remaining model parameters following the procedure in Appendix E. Table F.6 reports the estimated parameters and model fit.

Results. Figure F.14 reports the transition path for the amenity-spillover extension. Across the full transition, Tokyo's terminal population share slightly more than doubles, while Hokkaido's

[13] Fajgelbaum and Gaubert (2020) use the estimates from Diamond (2016) to estimate the following parameters: $\gamma_{UU} = -0.43$, $\gamma_{SU} = 0.18$, $\gamma_{SS} = -1.24$, $\gamma_{US} = 0.77$. Their definition of these terms implies that for a constant spillover elasticity to be consistent with these estimates, it must be between $\gamma_{UU} + \gamma_{SU} = -0.25$ and $\gamma_{SS} + \gamma_{US} = -0.47$.

Table F.6. Parameter estimates and model fit: amenity spillovers

Parameter	Value	Moment	Source	Model	Data
κ	$10^{6.45}$	Average out-migration rate	Japan 1990-2000	0.024	0.024
ϱ	1.31	Tokyo housing supply elasticity	Baum-Snow and Han (2024)	0.210	0.210
δ	0.08	Housing depreciation rate	Japan 1995-2024	0.058	0.058
ν	0.22	Construction factor income share	Japan 2000	0.085	0.085
z_Y/z_X	9.10	Construction empl. / income	Japan 2000	1.149	1.151
η	0.63	Housing supply vs. density slope	Baum-Snow and Han (2024)	-0.097	-0.097
α	0.34	Housing expenditure share	Piazzesi and Schneider (2016)	0.250	0.250
$\sum_i A_i$	14.20	Aggregate GDP per capita	Normalization	1.002	1.000

Notes: This table reproduces the results in Table 4 when estimating the model with the amenity spillover extension.

falls by 45%. These effects are roughly one-third to one-half as large as in our baseline model. In anticipation, the pattern in population shares is similar to the baseline model, with the magnitudes dampened by a similar amount. In contrast, the changes in house prices are qualitatively different than in the baseline model, with Hokkaido's house prices rising slightly in anticipation and falling by less along the full transition path. All of these changes are consistent with the negative amenity spillover dampening the concentration force, which in turn dampens the demand-side force that pushes down house prices in Hokkaido on a relative basis (see the discussion around (25)). While these results are specific to a particular choice of γ, we have experimented with other values and have in general found that our results are monotonic in the size of this spillover: more positive values, as Giannone et al. (2026) suggest for Japan, would amplify our results.

Figure F.14. Population shares and house prices during the transition: amenity spillovers

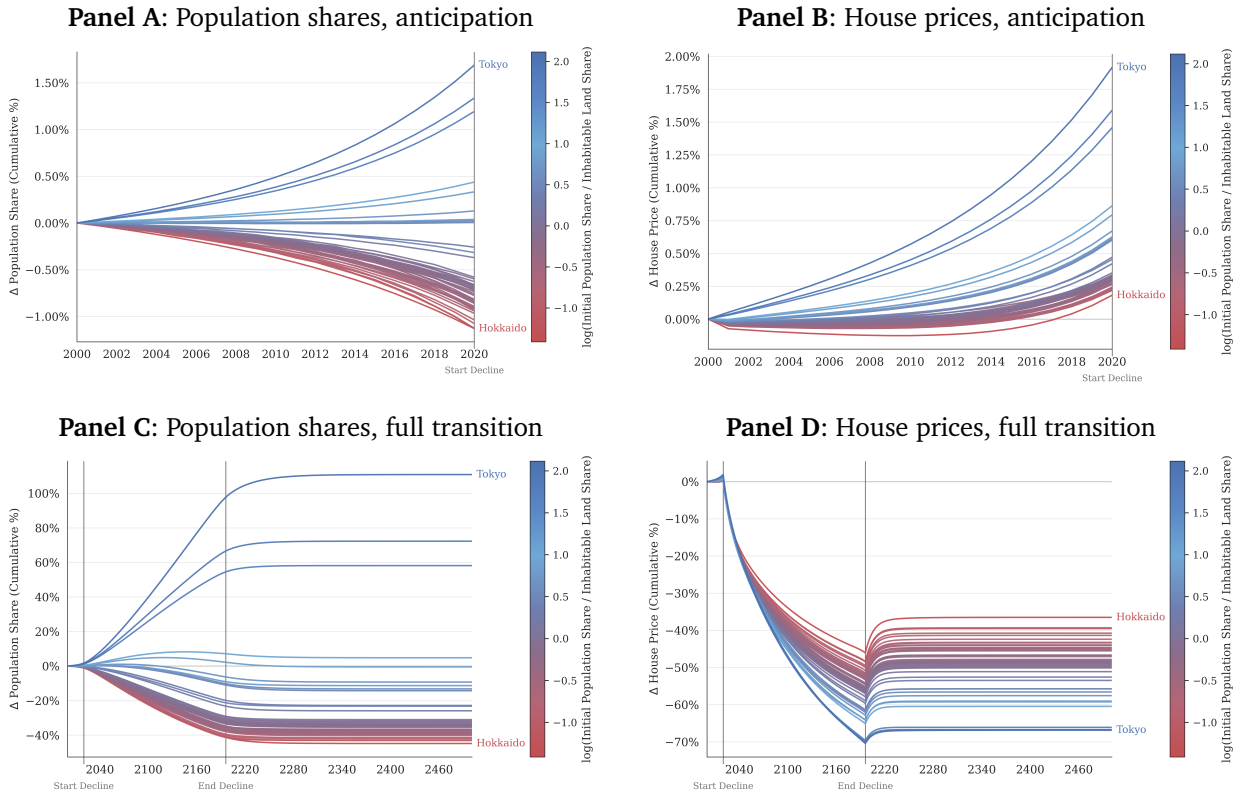

Notes: This figure reproduces Figure 8 in the model with the amenity spillover extension.

F.5 Extension: endogenous fertility

Here we consider what would happen in the model if fertility endogenously responded to housing supply, consistent with the evidence in Fazio et al. (2025). We focus on the implications of our key model assumption: CES housing supply with complements instead of Cobb-Douglas.

We adapt the economic setting (27)–(39) as follows. The utility function $U(c,h)$ is now also increasing in births b, becoming $U(c,h,b)$. The opportunity cost of child-rearing is the wage, so the budget constraint becomes $(1-b_{jot})w_{jot} \geq c_{jot} + p_{jt}h_{jot}$. This gives rise to indirect flow utility $U^*(w_{jot}, p_{jt})$ and Marshallian demand functions for goods consumption $c(w_{jot}, p_{jt})$, housing consumption $h(w_{jot}, p_{jt})$, and births $b_t(w_{jot}, p_{jt})$.[14] To directly capture the intuition that fertility responds to housing consumption, we assume that utility is such that the birth rate can be written in terms of just housing consumption, $b(w_{jot}, p_{jt}) = b(h_{jot})$. We give an example below of a simple setup that gives rise to this. Finally, we assume households in all locations die between periods at a mortality rate d. Under these assumptions, population dynamics (31)–(32) become

$$N_{jt} = \sum\nolimits_i N_{i,t-1}(1+g_{it})\mathbb{P}(j_{kt} = j \mid j_{k,t-1} = i), \tag{F.9}$$

$$g_{it} \equiv \sum\nolimits_o \frac{N_{io,t-1}}{N_{i,t-1}}(b_{t-1}(h_{io,t-1}) - d). \tag{F.10}$$

This then implies aggregate population growth

$$\frac{N_t}{N_{t-1}} = 1 + g_t = 1 + \sum_i \frac{N_{i,t-1}}{N_{t-1}} g_{it}. \tag{F.11}$$

If birth rates $b_t(h_{iot})$ are increasing in housing consumption, then, all else equal, an increase in per-capita housing supply will raise the aggregate population growth rate.

We are interested in the conditions under which this setting gives rise to a stable population level ($\lim_{t\to\infty} N_t = N^* \in (0,\infty)$), an "expanding cosmos" ($\lim_{t\to\infty} N_t = \infty$), or an "empty planet" ($\lim_{t\to\infty} N_t = 0$).[15] This will ultimately come down to two things: the responsiveness of per-capita housing supply $h_i = H_i/N_i$ to population level N_i, and the responsiveness of the fertility rate $b(h_i)$ to per-capita housing supply. Assume that the birth rate is strictly increasing in h, with bounds

$$\lim_{h\downarrow 0} b(h) = \underline{b} \quad \text{and} \quad \lim_{h\to\infty} b(h) = \bar{b}. \tag{F.12}$$

Of course, if $\bar{b} < d$, then we will always drift toward an empty planet; and if $\underline{b} > d$, we will always have an expanding cosmos. Assume then that $\underline{b} < d < \bar{b}$, so that it is in principle possible to have

[14] The time subscript on the birth rate allows for exogenous shifts in preferences that can initiate a decline in fertility.
[15] We borrow this terminology from Jones (2022).

either growth or decline. In this case, long run population dynamics will depend on the behavior of per-capita housing supply. For Cobb-Douglas housing supply and CES housing supply with complements, per-capita housing supply falls toward zero as $N_i \to \infty$. What is unclear is whether housing supply is infinite or finite on an empty planet. Under our assumption of CES housing supply (36), for example,

$$\lim_{N_{it} \downarrow 0} \frac{H(X_{it}, L_i)}{N_{it}} = \begin{cases} \infty & \text{if } \eta \geq 1, \\ (1-\nu)^{\frac{\eta}{\eta-1}} A_{iX} \hat{N}_{iXt}(0), & \text{if } \eta < 1, \end{cases} \tag{F.13}$$

where $\hat{N}_{iXt}(0) \in (0, 1]$ is the equilibrium employment share in construction. Thus, in a comparative static sense, local birth rates satisfy

$$\lim_{N_{it} \downarrow 0} b_{it} = \begin{cases} b(\infty) = \bar{b} & \text{if } \eta \geq 1, \\ b((1-\nu)^{\frac{\eta}{\eta-1}} A_{iX} \hat{N}_{iXt}(0)) < \bar{b}, & \text{if } \eta < 1. \end{cases} \tag{F.14}$$

Under Cobb-Douglas or substitutes ($\eta \geq 1$), the population eventually falls to a low enough level that housing supply becomes arbitrarily large, rising to the upper bound $\bar{b}$. Since $\bar{b} > d$, there exists a population level low enough that fertility rises above the death rate, inducing positive local population growth. Since local population $N_i \leq N$ everywhere, there must then also exist an $N^* > 0$ such that aggregate population growth $g_t > 0$ when $N_t < N^*$.

With complements ($\eta < 1$), the story is potentially very different. Housing supply per capita is bounded above, and hence fertility may never rise enough to stabilize the population. For example, suppose $b((1-\nu)^{\frac{\eta}{\eta-1}} A_{iX}) < d$ for all i. Then we will inevitably reach an empty planet. The reason this happens in the complements case is that, as population declines, so does the supply of labor available to build structures. Land per capita becomes abundant, but structures become a weak link everywhere. Having infinite land is of no value without any people to build houses on it.

This same logic applies to population dynamics in response to a fertility shock. Suppose that we start out in an initial steady state $N_0 = N^*$ with corresponding fixed housing supplies $\{H(X_j, L_j)\}$ and regional population growth rates $\{g_j\}$ and shares $\{N_j/N\}$.[16] Then, suppose there is a preference shock that reduces fertility, so that birth rates fall instantly everywhere. Aggregate population will begin to decline. As population declines, housing supply will tend to rise everywhere. Under Cobb-Douglas, population will fall until housing supply has risen enough to stabilize growth back to $g_t = 0$. Under CES with complements, population may fall indefinitely.

[16] In a steady state, locations with lower per-capita housing supply (like Tokyo) have $g_i < 0$, but this is offset exactly by above-average in-migration rates $\{\mathbb{P}(j_{kt} = i \mid j_{k,t-1} = j)\}_j$, consistent with fixed equilibrium population shares.

Figure F.15 illustrates these dynamics for Cobb-Douglas housing supply and CES with complements. Panel A shows the dynamics before a permanent negative fertility shock and Panel B shows the dynamics after. Before the shock (Panel A), both cases give rise to a positive steady state population level. This level is stable, in that population always drifts toward that level asymptotically. In Panel B, however, the fertility decline is large enough that, under CES with complements, the birth rate can never rise enough to stabilize population—we drift toward an empty planet.

Figure F.15. Endogenous fertility and population dynamics

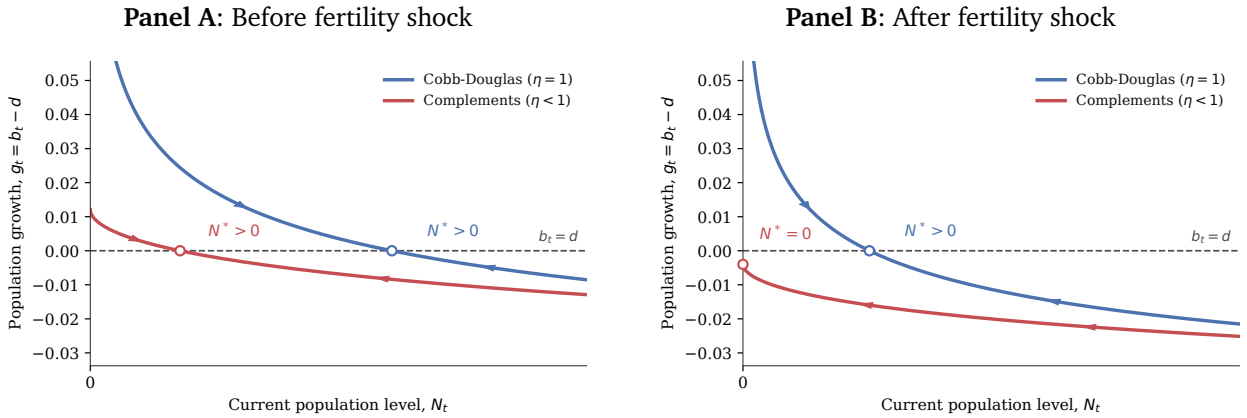

Notes: The figure plots the aggregate population growth rate implied by endogenous fertility under Cobb-Douglas housing supply and CES housing supply with complements. The horizontal dashed line marks zero population growth, $b_t = d$. Panel A uses a high fertility-preference shifter, while Panel B lowers this shifter enough that the CES-complements economy has no positive steady state population level.

In this particular example, we turn off migration costs ($\kappa_{ij} = 1$), housing durability ($\delta = 1$), and occupation shocks ($\varrho = 0$). We assume the direct utility function

$$U(c_{jt}, h_{jt}, b_{jt}) = c_{jt}^{1-\alpha} h_{jt}^{\alpha} b_{jt}^{\varepsilon_t^b(h_{jt})} \omega_t(h_{jt}), \tag{F.15}$$

where

$$\varepsilon_t^b(h_{jt}) \equiv \bar{\varepsilon}_t^b h_{jt}^{\zeta_b}, \tag{F.16}$$

and where

$$\omega_t(h) \equiv \frac{(1 + \varepsilon_t^b(h))^{1+\varepsilon_t^b(h)}}{\varepsilon_t^b(h)^{\varepsilon_t^b(h)}}$$

is a normalizing term. This term offsets the direct effect of $\varepsilon_t^b(h)$ on the marginal utility of housing, so the first-order conditions take the simple form

$$c_{jt} = (1 - \alpha)(1 - b_{jt}) w_{jt} \tag{F.17}$$

$$p_{jt} h_{jt} = \alpha (1 - b_{jt}) w_{jt} \tag{F.18}$$

$$b_{jt} = \frac{\varepsilon_t^b(h_{jt})}{1+\varepsilon_t^b(h_{jt})} = \frac{\bar{\varepsilon}_t^b h_{jt}^{\zeta_b}}{1+\bar{\varepsilon}_t^b h_{jt}^{\zeta_b}}. \tag{F.19}$$

If $\zeta_b > 0$, then fertility endogenously increases in housing consumption; if $\zeta_b = 0$, then fertility is exogenous. Figure F.15 assumes that the constant $\bar{\varepsilon}_t^b$ takes a high value in Panel A and a low value in Panel B. Both assume $\zeta_b > 0$. The remaining model parameters are set as in the main text.

F.6 Differences in optimal transfers between static and dynamic models

This section discusses the differences between the sign of the optimal transfers in our dynamic model and the optimal formula from Proposition 5 that applies in our static model. We begin by replicating the latter quantitatively by removing dynamics in the household and construction problems, taste shocks across occupations, and agglomeration, in which case the model becomes our static model. The first column of Table F.7 shows that the optimal τ_w is between -0.66 and -0.68, depending on the specification of housing supply and demand, which is close to the redistribution slope $-\theta/(\theta+1) \approx -0.67$ in (47). The second column of Table F.7 adds back agglomeration, which has minimal effects because the amount of agglomeration is small.[17] The third column introduces taste shocks across occupations. These shocks generate construction wage premia and make wage-indexed transfers redistribute across occupations within locations as well as across locations. With CES housing supply, within-location occupation wage differences are smaller in denser areas that tend to have higher wages, which causes τ_w to increase to around -0.56. In the fourth column, we add back dynamics in the construction sector, which reduces the demand for construction labor and construction wages. These effects differ across places with PIGL demand, which again interacts with the redistributive motive and causes τ_w to rise closer to zero.

The first four columns of Table F.7 show that occupation-specific taste shocks, depreciation, and agglomeration cannot explain the differences in the sign of τ_w between our static and dynamic models. The final column shows that the difference comes from moving costs. As discussed in Section 6.2, moving costs create differences in utility across places that allow the planner to improve welfare in steady state by moving people to higher-utility places.

[17] The optimal τ_w does not change exactly according to (47) because agglomeration is based on N_{iY} rather than N_i as in our static model. Therefore, the planner makes transfers to shift people from regions with high to low construction employment shares. If we change agglomeration to be based on N_i, the optimal τ_w increases, consistent with (47).

Table F.7. Optimal transfer policies across model specifications

Housing		Optimal τ_w				
Supply	Demand	(1)	(2)	(3)	(4)	(5)
Initial Steady State						
CD	CD	-0.68	-0.68	-0.68	-0.54	0.21
CES	CD	-0.67	-0.68	-0.56	-0.41	0.17
CD	PIGL	-0.68	-0.69	-0.68	-0.66	0.17
CES	PIGL	-0.67	-0.69	-0.56	-0.54	0.23
Terminal Steady State						
CD	CD	-0.68	-0.68	-0.68	-0.54	0.21
CES	CD	-0.66	-0.68	-0.56	-0.28	0.46
CD	PIGL	-0.68	-0.67	-0.69	-0.73	0.41
CES	PIGL	-0.67	-0.68	-0.57	-0.54	0.52
Parameters						
σ		0	0.05	0.05	0.05	0.05
ϱ		0	0	1.20	1.20	1.20
δ		1	1	1	0.08	0.08
κ		1	1	1	1	$10^{6.45}$

Notes: The table reports the planner's optimal value of τ_w in the post-transfer expenditure rule $x_{jot} = (1 + \tau_w)w_{jot} + \tau_0$, with τ_0 chosen to balance the government budget, solved separately in the initial and terminal steady states. The first two columns indicate the specifications of housing supply and demand, and Columns (1)–(5) vary the parameters listed in the bottom panel. The parameter values in the bottom panel correspond to the baseline estimates with CES supply and PIGL demand, reported in Table 4; for the alternative models, any parameter not fixed externally uses the corresponding estimate reported in Table F.1, Table F.2, and Table F.3.