Lopsided Cities: The costs of lopsided jobs distributions, and place-based policy responses

Xavier Tierney

Draft as of OCT 2026

Latest version at `https://xavier-tierney.github.io/files/Tierney_JMP.pdf`

Abstract

Many cities have lopsided employment distributions: the center of business is tied to one side of the city, while residences sprawl out in the other direction. This means average distances to employment are further than they would be in a symmetric city – increasing commute times. In this paper, I investigate the costs of such lopsidedness, its impact on inequality and possible policy responses. Calibrating a quantitative urban model to Sydney, I show that a counterfactual city with a more centralised employment hub has 11% higher output and 2.5% higher worker welfare, with higher gains for graduates. I then investigate the effects of a government policy to try to directly alter the spatial distribution of employment: the 2018-2022 shift of over 10,000 government jobs to a more central location in Sydney. Using a synthetic control method, I find that land values surrounding the new offices increased substantially compared to comparable locations. I also find that local approvals for new private office constructions increased by over 1 billion AUD, consistent with space for 7,000 new private workers, though there is insufficient post period data to test whether they have been realised. Using the model to consider general equilibrium effects of the public sector jobs move, I show these significant local effects are consistent with small effects for the city as a whole: total output falls slightly along with welfare of high skill workers due to decentralisation out of the CBD, but welfare of low skill workers marginally increases.

1 Introduction

In most cities, a substantial portion of employment and output is concentrated in a relatively small, densely occupied area, often centered around a single central buisiness district (CBD). Because these employment clusters are sustained by self reinforcing agglomeration externalities, their location is highly persistent over time. If geographic barriers or planning choices lead to asymmetric city sprawl, the center of business can therefore be left stranded to one side of the city while residents spread in the other direction. This lopsidedness will mean

that average distances from homes to the highest concentration of employment will be further than in a symmetric city.

Sydney's colonial settlment is an archetype of the pattern described above: British settlements were first built around the harbor of Port Jackson, and as the city has expanded it has run into the ocean to the east and mountains to the north and south, forcing it to sprawl west. The central business district, meanwhile, has remained firmly rooted in the east around the original colonial buildings on Sydney Cove. The extent of this westward sprawl is evident in the fact that the median place of residence is 14 km West of the CBD [1]. And as a result, the average resident is 25 % further from the CBD than from the most centrally situated alternative [2]. Given that the CBD contains 19 % of the cities 2.3 million jobs, and 43 % of its high paying jobs [3], it is unsurprising that residents' access to jobs and incomes are likewise geographically lopsided: for the 50 % of residents living to the west of the median residence, incomes are 29 % lower, and there are 617,000 fewer jobs located in that half of the city.

In General equilibrium, several forces could moderate the impact of a lopsided employment distribution on realised commutes, output and welfare. Individuals can reduce commutes by sorting into areas closer to where they work: if the high skill individuals most demanded by the firms predominant in the CBD can outbid others for homes in the surrounding suburbs, the cost in terms of commute times and output could be minimised. This same sorting also effects the inequality impacts of the jobs distribution: on the one hand, low income workers are most likely to be pushed to the far periphery as they least able to pay for access to the CBD (and indeed the median non graduate is 17km West of the CBD, vs 9km for gradutes); on the other, poor access to the CBD hurts gaduate workers more because CBD firms are the most graduate intensive (the CBD's workforce is 62% graduates, vs 33% for areas more than 10km from the CBD). Additionally, residents and firms could respond to the poor access to employment and agglomeration externalities on the far side of the city by densifying more intensely around the employment center. Indeed, by imposing a higher cost to sprawl, geographic constraints could incidentally internalise agglomeration benefits and thereby increase output.

To understand the costs of a lopsided employment distribution, and the channels through which they operate, I therefore turn to a spatial equilibrium model. The baseline model is similar to Ahlfeldt et al (2015) [3], in which the natural sorting of workers into homes closer to their place of work is captured by the fact that when comparing residence-workplace location pairs they trade off commute costs against rental prices, amenities and wages. To capture the substantial sorting of types across the city, I also introduce two distinct worker types, split by level of education (as does eg Tsivanidis (2026) [16]). I allow the relative graduate intensity of production to vary across space, the attractiveness

[1] calculated as the median latitude and longitude of residents in the Sydney SUA by place of residence by SA1 in the 2021 Census

[2] the geometric median, a few hundred metres from the marginal median described above

[3] 43 % of those in the highest income band in the census data, the top 8.4% of workers with positive stated income

of locations to the two types to vary across space, and the low skill types to be more sensitive to rental prices as a result of higher housing expenditure shares.

To benchmark the costs of Sydney's lopsided employment distribution, I compare it to a counterfactual city with a more central CBD. First, I calibrate a baseline of the model to census and land value data from Sydney. I then shift the fundamental amenity and productivity of the $9km^2$ around the CBD from its current location to a more central location, 20km west in Parramatta. This more central location has greater access to both workers and adjacent developable land, and as a result employment concentrates more densly here than around the baseline harborside CBD location: the difference peaks in the densest $13km^2$) of the counterfactual city, which contains 294,000 more workers than are employed the densest $13km^2$of the original city. With more firms able to locate in the highest productivity areas, and further productivity gains from increased agglomeration spillovers due to the higher density itself, overall output increases by 11.0 % relative to the baseline cit. With the concentration of employment in a more central location, average commute times fall by 2.1%, and along with higher wages this leads to worker welfare increasing by by 2.5 %.[4] This worker welfare is before accounting for any profits they receive as landlords/homeowners, and landlord profits rise by the same as output– 11.0 %. As the CBD production technology is more graduate intensive, the greater CBD access and hence centralisation in the counterfactual city helps graduates more than non graduates: worker-renter welfare rises by 3.7% for graduates but falls by 0.1% for non graduates. This runs counter the naive intuition that non-graduates might be hurt more by the CBD's off center location because they are pushed to a further periphery than in a more central location, and shows the importance of accounting for exactly which types of jobs have poor access in the city.

This stylised counterfactual suggests that there could be significant gains from shifting employment to a more central location. One particularly direct way for policymakers to do so is to move the jobs directly under their control: civil service jobs. In the recent 'decade of decentralisation' policy, the NSW government did just this, reducing government CBD office holdings by over 100,000 sqm and rehousing over 10,000 workers in new offices 20 km West of the CBD in Parramatta. This policy was partly aimed at simply reducing exposure to expensive government rents and placing civil servants closer to the people they serve. But statements in planning documents and press releases suggest it was at least partly motivated by the desire to make jobs easier to access both for civil servants themselves and for other Western Sydney residents, by encouraging private buisinesses to locate around the new government offices and promote Parramatta as a viable 'second CBD'.

To understand the effects of the policy in reduced form, I run a synthetic control analysis [1] using a panel of land values for all properties in greater Sydney from 2012-present. I take a pool of comparable commercial centers,

[4]Holding residence amenities fixed at baseline rather than letting them respond to the re-sorted population, the welfare gain is 1.3% and the output gain 7.0%.

and identify the synthetic control as the weighted combination of these centers that best match the pre-treatment trajectory of outcomes in Parramatta. Comparing the post treatment land values for Parramatta and this synthetic control, I find that Parramatta had substantially higher post treatment landvalues, with an estimated treatment effect of +$1,700 per square metre in my headline specification and larger post treatment effect than all 15 of the eligible donor pool members. Additionally, running the same synthetic control procedure on office approvals, i find Parramatta had AUD$1.05 billion[5] more private office construction approved than its synthetic control and a larger post treatment effect than all 10 of the eligible donor pool members. Comparing this to the floorspace/construction cost of the government controlled towers, and the floorspace per private employment in the CBD floorspace use survey, I find this extra construction could house around 7,000 new private office jobs. The latest census employment data from 2021 is confounded by covid and taken before a large portion of the new civil servant jobs were fully settled or the new private office construction was completed, so it is difficult to identify whether this extra office space has been realised as extra private sector jobs. Employment up to 2021 appears to suggest no substantial impact on non-public service jobs to that point, with possible increases in professional, scientific and technical services.

These reduced form results give us a sense of the local impacts of the policy, but cannot capture general equilbrium effects on the rest of the city or overall welfare. I therefore adapt my quantitative urban model to include a public sector so that i can study the impacts of the policy in the model equilibrium. The public sector is seperate from the private sectors, and is controlled by a policy maker that dictates where to rent offices for its workers. The key model parameter is the spillover from public to private sector firms: if these spillovers are high, the government can alter pull private sector jobs across the city just by moving its own workers.

To understand the impact of the policy I compare the calibrated equilibrium to the equilirbium in which 14,000 workers are moved west. Because the policy is relatively modest in comparison to the previous extreme CBD shifting counterfactual (exogenously moving only 14,000 workers compared to the 500,000 that work in the CBD), the citywide effects of the office moves are correspondingly modest: if we assume public-private spillovers are the same as for private workers (at an employment density-productivity elasticity of 7%), output goes down by 0.12%, while graduate workers welfare goes down by 0.058% and non graduates welfare goes up by 0.040%. The sign of the results turn on the spillover from public to private sector firms. When spillovers are high, moving public sector jobs leads to higher employment in their new location and lower employment around the CBD. Because the CBD is the highest productivity location, this decentralisation leads overall output to fall. Because the CBD is closer to where graduates live and uses a more graduate intensive production technology, graduates lose from this dispersion; non graduates live closer to the new hub and the production technology there is less graduate intensive, so non gradu-

[5]in FY2025-26 construction prices, or 946m in 2026 Sydney CPI adjusted value

ates gain from this dispersion. When public-private spillovers are low, however, moving public sector jobs out of the CBD just frees up space for private sector workers there and reduces spaces where they relocate to, so the effects switch: workers centralise more in the high productivity, graduate intensive CBD, and output and graduate welfare go up while non-graduates lose.

The local productivity increase needed for the model to explain the reduced form land values is 5.6%; to explain the increase in employment needed to fill the extra space for private offices implied by the reduced form approvals estimates it would need to increase by 2.6%. The model requires a very large value of public to private spillovers to explain these productivity estimates, which would imply gains to low type workers at the expense of high type workers and welfare.

Comparing the results of the CBD shift counterfactual to the public office sector counterfactual reveals some key lessons for policymakers. First, the costs of the CBD's lopsided location are substantial. However, even substantial policies like moving a large portion of civil service workers west cannot recoup much of those costs, because they are too small to shift the CBD and materially impact the overall distribution of jobs. Indeed, by decentralising jobs away from the most productive location they are likely to actually reduce overall output. Nonetheless, the existence of a neglected periphery with low employment access does present an opportunity to reduce inequality through place based policies to move employment to those areas. This opportunity exists despite the fact that the CBD's lopsided location actually hurts high type workers more than low type workers, and the explanation is the same for both phenomena. The CBD is a high productivity, graduate intensive location, hence moving it wholesale to a more accessible location increases output and graduate wages. Conversely, small scale policies that only decentralise employment to less graduate intensive areas will reduce output and graduate welfare, while improving non graduate welfare.

1.1 Literature review

The primary existing study of the costs of bad shape in the economics literature is Harari 2020 [11]. Harari uses growth running into geographic constraints as an instrument for non-compact city shape, and finds that a standard deviation worse shape is associated with 3% smaller population growth. She also finds wages in cities with bad shape appear if anything higher (though the effect is not statistically significant)– which she interprets through a Rosen-Roback model as being consistent with compensating differentials for worse travel times in the city.

Because I have detailed within city data on where workers live and work, I am able to consider the particular issue of lopsided jobs distribtions within a fixed urban boundary, calibrate a quantitative urban model to understand the channels through which it impacts the city, and consider a real world policy designed to counteract it. In an appendix, I also consider the costs of non-compact urban boundaries through the lense of the model, providing corroberation to Harari's findings that urban boundaries are costly through a different, structural method,

while also allowing me to probe specific aspects of urban boundaries and the channels through which they operate.

There is a also more broad literature on the costs of sprawl, see eg. [10]. Perhaps closest to this paper, Koster 2024 [14] uses a quantitative urban model to assess the impacts of London's greenbelt, finding that the greenbelt's removal of developable space increases rents, but the estimated increase in amenities suggests the greenbelt is overall welfare improving. Instead of looking at the impacts of external boundaries to expansion, this paper considers the complementary issue of how jobs are distributed within those boundaries

My analysis of the movement of jobs to Parramatta relates to a growing literature on the local effects of public service job relocation. In a review of the literature, Freitas [9] finds that effects differ across studies but concludes that public-job relocations produce an average short-run public-to-private multiplier of approximately 0.7, driven mainly by non-traded activity, while their long-run effects on tradable employment remain uncertain. Studies of actual office relocations generally find positive but highly localized effects. Faggio [7] shows that the UK Lyons Review increased private-service employment close to receiving sites, but also drew firms inward from surrounding neighbourhoods. Faggio et al. [8] likewise find that the relocation of German government functions increased private employment primarily in services and within a few hundred metres of the receiving offices. Becker, Heblich and Sturm [5] find that Bonn's designation as the West German capital created more than one additional non-industrial private job per public job while displacing some industrial activity, without substantially increasing overall local productivity. The BBC's move to Salford changed creative-industry employment and firm composition without significantly increasing total employment [15], while South Korea's public-agency relocations generated limited employment beyond the transferred public workers themselves [13].

This literature has however focused on cross city moves, and devoted little attention to moves of public sector workers within a city. However, many governments have moved public sector jobs within cities, including in Tokyo, Barcelona, and Beunos Aries. And it is easy to see why it may be attractive to local policymakers. Local policymakers are acutely aware of spatial inequality in employment access in their city but may not have the power to redistribute large sums through income taxes, or have the budgets for large infrastructure projects. However,they can move the jobs directly under their control– their civil servants– and doing so has the political benefit of appearing to very directly tackle spatial inequality in jobs access. Sydney is a particularly attractive setting for such a policy, because the lopsided jobs distribution means there is a large part of the city with surplus local labour supply. The indicative reduced form findings of the impacts of the Sydney jobs move on private office approvals are consistent with the broader literature's findings of approximatly 0.7 extra local jobs for each private sector job. The model's general equilirbium interpretation of the policy show the tradeoff inherent in the policy: decentralising workers can modestly reduce inequality, but at the cost of lower output and reduced graduate welfare, even in a setting in which the existing CBD has as

poor access as Sydney's.

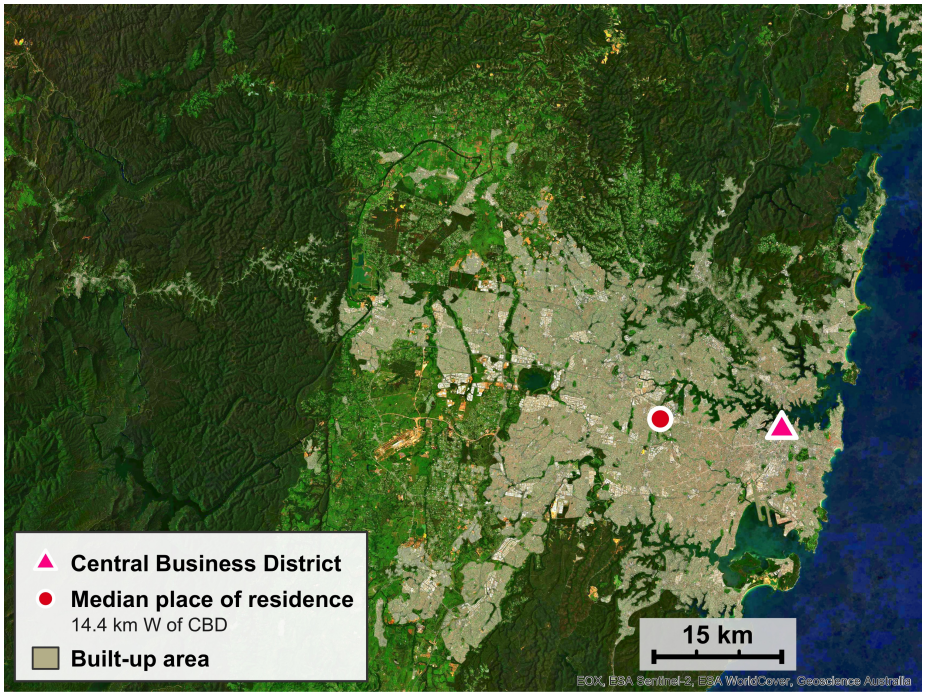


Figure 1: Sydney from satellite imagery, shaded by relief. The imagery is lightened where land cover is built-up. The city occupies a coastal plain bounded to the west by the Blue Mountains escarpment, to the north by the Hawkesbury river system and to the east by the ocean; the harbour and its estuaries cut deeply into the built-up area. The hot-pink triangle marks Martin Place, the CBD anchor used throughout, and the red circle the resident-weighted median place of residence, 14.4 km west of it (2021 census, SA1 place of residence). Imagery: EOxCloudless (`https://cloudless.eox.at`) by EOX IT Services GmbH (contains modified Copernicus Sentinel data 2024), CC BY-NC-SA 4.0, colour-adjusted. Land cover: © ESA WorldCover project 2021, contains modified Copernicus Sentinel data (2021) processed by ESA WorldCover consortium, CC BY 4.0. Relief: AWS Terrain Tiles; Australia terrain data © Commonwealth of Australia (Geoscience Australia) 2017, CC BY 4.0; SRTM and GMTED2010 courtesy of the U.S. Geological Survey; ETOPO1 courtesy of NOAA.

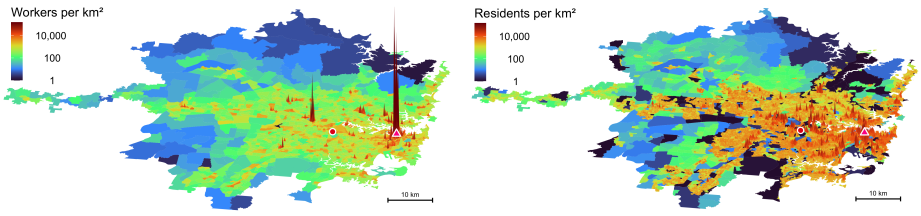


Figure 2: Employment density by destination zone (left) and resident density by SA1 (right), Sydney SUA, 2021 census. Oblique view: colour and spike height both encode density (persons per km^2 of unit area; log colour scale; spike height linear in density). The two panels share a single colour scale and a single spike-height scale, so they are directly comparable: employment is concentrated in a single dominant CBD cluster in the city's east — its peak density is seven times the residential maximum — while residence is spread across the whole urban footprint. The hot-pink triangle marks Martin Place, the CBD anchor point used throughout the paper; the red circle marks the resident-weighted median place of residence, 14 km west of it.

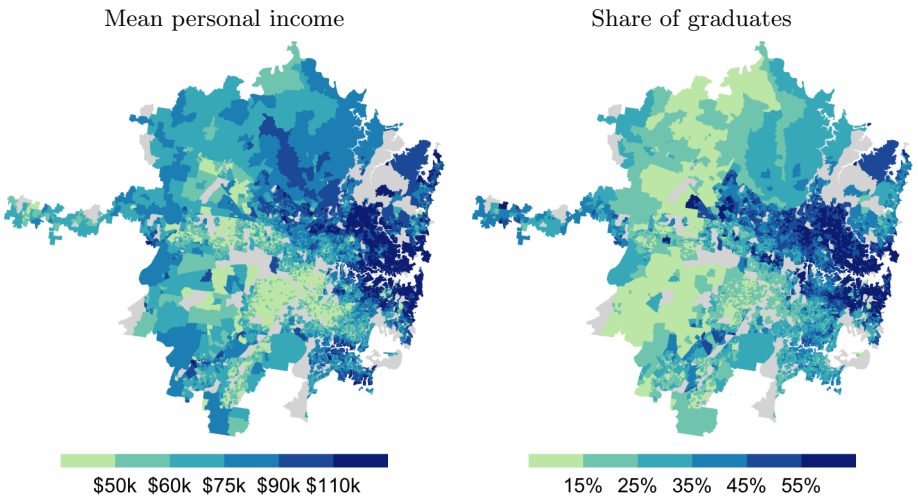


Figure 3: Income and graduates by place of residence, SA1s of the Sydney SUA, 2021 census. Left: mean annual personal income of residents with a positive stated income (banded census incomes at band midpoints). Right: share of residents with a bachelor's degree or higher, among those whose highest attainment is classified. Grey: fewer than 20 people in the measure's base.

2 Static model

To benchmark the costs of Sydney's lopsided employment, I first run a counterfactual with a static version of the model, which is identical to the long run equilibrium of the dynamic model (described later) under appropriate parameters.

2.1 Overview

The static model is similar to the standard quantitative urban model from Ahlfeldt et al, but I split workers by level of education to capture their spatial sorting. In the baseline model, the types are high education (Bachelor's degree or above) and low education (below degree level), which divides Sydney's employed population nearly in half. Workers choose where to live and work by trading off wages where they work, rental prices and type specific amenities where they live, and commute times. Low type workers have higher expenditure share on floorspace, leading them to sort more toward places with lower rent. Firms are split into two types: one that hires low type workers only, the other that hires high type workers only, and relative labor demand for the two types of workers in a given place will be determined by the location's relative fundamental productivity for high vs low type firms. A construction sector responds to higher firm and resident willingness to pay rent by providing more floorspace. In equilibrium, prices adjust so that demand and supply for workers and floorspace are matched.

2.2 Workers

The utility of worker o of type t when living in i and working in j is a Cobb douglas function of goods and floorspace consumption, leading to indirect utility of the form:

$$u_{ijo}^{t} = \frac{B_i^t w_j^t z_{ijo}^t}{Q_i^{1-\beta^t} d_{ij}^t} \tag{1}$$

Where

- w_j^t is the wage for type t in location j
- Q_i is the rental price of floorspace in location i
- The expenditure share on housing is $1 - \beta^t$, and is smaller for low types, so they have greater expenditure share on housing. This means low type households will perceive a higher cost from living in areas with higher rent, and will therefore lead them to sort more strongly to areas with lower rent
- B_i^t is the 'amenity' of location i for type t. This captures anything that makes a location more attractive to a given type that is not captured in rental prices and access to jobs, which could include standard amenities like views of the harbor, but also type specific local non tradeable goods

(high vs low end restaurants) or the presence of a community of other individuals of the same type.

- z_{ijo}^t is o's idiosyncratic utility of living in i and working in j
- d_{ij}^t is the cost of commuting from i to j, which takes form $d_{ij} = e^{\kappa^t \tau_{ij}}$ where τ_{ij} is the travel time from i to j

An individual's idiosyncratic preference shock z_{ijo} comes from a random draw from a frechet distribution with frechet shape parameter ε, common to both types (a higher ε means preferences are more similar across people). Having observed their preference draw, consumers choose the utility maximizing commute pair, and if we aggregate over the continuum of all workers' draws we find (via the properties of the maximum of the Frechet distribution) that the fraction of all workers who live in i and commute to j is given by:

$$\pi_{ij}^t = \frac{T_i E_j (\frac{B_i^t w_j^t}{d_{ij}^t Q_i^{1-\beta^t}})^\varepsilon}{\sum_{r,s} T_r E_s (\frac{B_r^t w_s^t}{d_{rs}^t Q_r^{1-\beta^t}})^\varepsilon} \tag{2}$$

which we'll label:

$$\pi_{ij}^t = \frac{\Phi_{ij}^t}{\Phi^t} \tag{3}$$

2.3 Production

Firms produce a single, perfectly tradable good in perfect competition with constant returns to scale. A representative firm in each location j hires both high and low type workers, and combines them with commercial floorspace. Its production function is Cobb–Douglas in labour and floorspace, where labour is a CES bundle of the two types:

$$Y_j = A_j (L_{Fj})^\alpha (H_{Fj})^{1-\alpha}, \qquad L_{Fj} = \left[(L_{Fj}^l)^{\frac{\sigma-1}{\sigma}} + \mu_j^h \, (L_{Fj}^h)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \tag{4}$$

where

- A_j is the productivity of location j, a single Hicks-neutral productivity common to both types
- μ_j^h is the graduate intensity of production in j: the weight on high type workers in the labour bundle, relative to low type workers, whose weight μ_j^l is normalised to one. One weight per location is a free normalisation, since scaling both weights by a common factor only rescales L_{Fj}, which A_j absorbs[6]

[6] An observed zero employment of a type in a location can only be rationalised by a zero weight on that type there. In locations where only high type workers are employed, the normalisation is placed on μ_j^h instead.

- $\sigma > 1$ is the elasticity of substitution between high and low type workers

Differences in graduate intensity across locations, along with differences in each location's access to workers of each type, cause the two types to sort differently across workplaces.

Profits are:

$$\pi_j = A_j (L_{Fj})^{\alpha} (H_{Fj})^{1-\alpha} - \sum_t w_j^t L_{Fj}^t - Q_j H_{Fj} \tag{5}$$

The firm's problem splits in two. First, minimising the cost of the labour bundle over the two types gives its unit cost w_{Fj}, a CES wage index, and the relative demand for the two types:

$$w_{Fj} = \Big[(w_j^l)^{1-\sigma} + (\mu_j^h)^{\sigma} (w_j^h)^{1-\sigma} \Big]^{\frac{1}{1-\sigma}}, \qquad \frac{L_{Fj}^h}{L_{Fj}^l} = (\mu_j^h)^{\sigma} \Big(\frac{w_j^h}{w_j^l} \Big)^{-\sigma} \tag{6}$$

with $w_{Fj} L_{Fj} = \sum_t w_j^t L_{Fj}^t$, so the cost of the bundle is the wage bill. Second, given the bundle's price w_{Fj}, the firm's choice of labour and floorspace is the one-type Cobb–Douglas problem with w_{Fj} in place of the wage. Profit maximisation with respect to floorspace gives commercial floorspace demand:

$$H_{Fj} = L_{Fj} \Big((1-\alpha) \frac{A_j}{Q_j} \Big)^{1/\alpha} = \frac{1-\alpha}{\alpha} \frac{\sum_t w_j^t L_{Fj}^t}{Q_j} \tag{7}$$

Given perfect competition there are zero profits in equilibrium, which pins down the price of the labour bundle at j as a function of local productivity and rents:

$$w_{Fj} = \alpha (1-\alpha)^{\frac{1-\alpha}{\alpha}} A_j^{\frac{1}{\alpha}} Q_j^{-\frac{1-\alpha}{\alpha}} \tag{8}$$

Each type's wage offer is then its marginal product within the bundle:

$$w_j^t = \mu_j^t \Big(\frac{L_{Fj}^t}{L_{Fj}} \Big)^{-\frac{1}{\sigma}} w_{Fj}, \qquad \frac{w_j^h}{w_j^l} = \mu_j^h \Big(\frac{L_{Fj}^h}{L_{Fj}^l} \Big)^{-\frac{1}{\sigma}} \tag{9}$$

The bundle is still demanded perfectly elastically at w_{Fj}, so total employment in a location is set by worker supply, but its composition now feeds back into wages. A type that is scarce in a location earns a premium there, and, at given rents, more high type workers in j raise the low type wage in j: $\partial \ln w_j^l / \partial \ln L_{Fj}^h = s_j^h / \sigma > 0$, where s_j^h is the high type share of the wage bill in j.

As $\sigma \to \infty$ the two types become perfect substitutes in efficiency units, and $w_j^h / w_j^l = \mu_j^h$ whatever the quantities employed. The model is then identical to one with separate high and low type firms, each with a Cobb–Douglas production function in its own workers and floorspace, with productivities $A_j^l = A_j$ and $A_j^h = (\mu_j^h)^{\alpha} A_j$.

2.4 Floorspace

I assume that floorspace is perfectly substitutable between commercial and residential uses and has a common price Q_i. Floorspace supply in a location is a constant elasticity function of rent:

$$H_i = K_i \psi_i Q_i^{\zeta} \tag{10}$$

where

- ζ is the elasticity of supply
- ψ_i is a location specific cost shifter, which could include costs due to zoning regulation or topography
- K_i is the land available for development in i

This could be microfounded by landowners taking the local floorspace rental price as given and construct floorspace according to a Cobb-Douglas functino of land and capital, as in Ahlfeldt 2015. It will be discussed in more detail in the dynamic model section.

2.5 Agglomeration spillovers

In line with the long literature on agglomeration spillovers, the productivity and amenity in a location are affected by the number of workers of each type in a location. Future drafts will attempt to estimate the particular values of these spillovers as a function of types, but for now I assume spillovers are type agnostic and take the same form as Ahlfeldt et al 2015. Denoting the fundamental productivity and amenity by a_j and b_i, the total productivity A_j and total amenity B_i depend on total employment $L_{Fs} = \sum_t L_{Fs}^t$ and total residents $L_{Rs} = \sum_t L_{Rs}^t$ per unit of land:

$$A_j = a_j (\sum_s e^{-\delta \tau_{js}} L_{Fs}/K_s)^{\lambda} \tag{E8}$$

$$B_i^t = b_i^t (\sum_s e^{-\rho \tau_{is}} L_{Rs}/K_s)^{\eta} \tag{E9}$$

2.6 Equilibrium

There are two worker types, $t \in \{h, l\}$. Given parameters

$$\Theta = \{\alpha, \sigma, \beta^h, \beta^l, \varepsilon, \kappa^h, \kappa^l, \zeta, \lambda, \delta, \eta, \rho\},$$

exogenous fundamentals

$$\{a_i, \mu_i^h, b_i^t, T_i, E_i, \psi_i, K_i, \tau_{ij}\},$$

and fixed total populations of each type L^h_{tot} and L^l_{tot}, the equilibrium is the vector

$$\{Q_i, A_i, H_{Fi}, w^t_i, B^t_i, H^t_{Ri}, L^t_{Fi}, L^t_{Ri}, \pi^t_{ij}\}_{t=h,l}$$

which satisfies the following equations for each type t. The two types interact through the common floorspace price Q_i in (E7); through production, since each type's wage in (E6) depends on the mix of the two types employed in the same location; and through the spillovers in (E8) and (E9), which depend on total employment and total residents.

Commute pair shares, where $d^t_{ij} = e^{\kappa^t \tau_{ij}}$:

$$\pi^t_{ij} = \frac{T_i E_j (\frac{B^t_i w^t_j}{d^t_{ij} Q_i^{1-\beta^t}})^\varepsilon}{\sum_{r,s} T_r E_s (\frac{B^t_r w^t_s}{d^t_{rs} Q_r^{1-\beta^t}})^\varepsilon} \tag{E1}$$

Residential population of type t in location i is

$$L^t_{Ri} = L^t_{tot} \sum_j \pi^t_{ij} \tag{E2}$$

Workplace employment of type t in location j is

$$L^t_{Fj} = L^t_{tot} \sum_i \pi^t_{ij} \tag{E3}$$

Residential floorspace demand, where $\pi^t_{j|i} = \pi^t_{ij} / \sum_s \pi^t_{is}$ is the probability that a type t resident of i works in j:

$$H^t_{Ri} = (1-\beta^t) \frac{\sum_j \pi^t_{j|i} w^t_j L^t_{Ri}}{Q_i} \tag{E4}$$

Commercial floorspace demand, where $L_{Fi} = \left[(L^l_{Fi})^{\frac{\sigma-1}{\sigma}} + \mu^h_i (L^h_{Fi})^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}$ is the CES labour bundle (4):

$$H_{Fi} = L_{Fi} \Big((1-\alpha) \frac{A_i}{Q_i} \Big)^{1/\alpha} = \frac{1-\alpha}{\alpha} \frac{\sum_t w^t_i L^t_{Fi}}{Q_i} \tag{E5}$$

Zero profits and wages, where $\mu^l_j \equiv 1$:

$$w^t_j = \mu^t_j \Big(\frac{L^t_{Fj}}{L_{Fj}} \Big)^{-1/\sigma} \alpha (1-\alpha)^{(1-\alpha)/\alpha} A_j^{1/\alpha} Q_j^{-(1-\alpha)/\alpha} \tag{E6}$$

Floorspace market clearing:

$$\sum_t H^t_{Ri} + H_{Fi} = K_i \psi_i Q_i^\zeta \tag{E7}$$

Productivity and amenity spillovers, which depend on total employment $L_{Fs} = \sum_t L_{Fs}^t$ and total residents $L_{Rs} = \sum_t L_{Rs}^t$ per unit of land:

$$A_j = a_j(\sum_s e^{-\delta \tau_{js}} L_{Fs}/K_s)^\lambda \tag{E8}$$

$$B_i^t = b_i^t(\sum_s e^{-\rho \tau_{is}} L_{Rs}/K_s)^\eta \tag{E9}$$

With $\lambda = \eta = 0$ there are no spillovers, and $A_j = a_j$, $B_i^t = b_i^t$.

2.7 Utility

Ex ante expected utility of a type t worker is

$$E[U^t] = \Gamma(1 - \frac{1}{\varepsilon})\,\bar{U}^t, \qquad \bar{U}^t \equiv (\Phi^t)^{1/\varepsilon} = (\sum_r \sum_s \Phi_{rs}^t)^{1/\varepsilon} \tag{11}$$

where $\Gamma(\cdot)$ is the gamma function and Φ_{rs}^t is the numerator of the commute share (E1) for the pair (r, s). The properties of the Fréchet distribution mean that, within a type, the distribution of utility is the same in every residence–workplace pair: higher wages or lower rents in a pair draw in workers with weaker idiosyncratic preferences for it, until the two exactly offset. Welfare inequality measures in the model are therefore between type measures, not between places within a type.

I measure each type's welfare change in a counterfactual as the proportional change in its expected utility:

$$\frac{\bar{U}^{t\prime}}{\bar{U}^t} = (\frac{\Phi^{t\prime}}{\Phi^t})^{1/\varepsilon} \tag{12}$$

where primes denote the counterfactual equilibrium. These changes can be compared across types and across counterfactuals. The level of $\bar{U}^h$ relative to $\bar{U}^l$, however, is not identified: scaling all of a type's amenities up by the same factor leaves its residential choices unchanged, so the calibration recovers each type's amenities only up to a type-specific scale.

3 Calibration: simple one type

3.1 Overview

I calibrate the model's parameters using data from Sydney. The process finds the parameters such that the data are consistent with the model in equilirbium. The parameters are exactly identified by the targeted moments (i.e. there is a one-to-one match), so that solving for the equilibrium using the calibrated parameters will exactly reproduce the targeted moments in the observed data. With the baseline parameters in hand, I can later find counterfactual model

equilibria with certain characteristics of the city changed but others held fixed. The calibration closely follows Ahlfeldt et al.

A number of parameters are taken from the literature. The agglomeration parameters are matched to Ahlfeldt et al 2015. The share of firm expenditure on floorspace is $(1 - \alpha) = 0.2$ following Valentinyi and Herrendorf 2008, as in Ahlfeldt et al 2015. Floorspace construction elasticity of $\zeta = 1.86$ is taken to match a simple model with land use share in construction of 0.35 following the housing production function estimates of Combes, Duranton and Gobillon 2021 [6], whose capital elasticity of 0.65 under constant returns Cobb–Douglas in land and capital implies $\zeta = 0.65/0.35$; this sits close to the 1.83 50 year response to rateable values calibrated for London by Heblich, Redding and Sturm 2020 [12]. This value also results in the observed log correlation of land values and rental prices across space closely matching the elasticity implied by my model, as described below. In robustness checks I lower it to 0.5 following Baum-Snow and Han 2024 [4] (also the value used for my short run response in my dynamic model), and raise it to 3.0, the value implied by the Cobb–Douglas construction function of Ahlfeldt et al 2015 [3]. Note that the static model is designed to match the long run equilibria of my urban model, and is being used in the next section to test counterfactual long run equilibria of a different city structure, which is why I base the headline elasticity on the technology function of Combes et al (2021)/ the 50 year response of Heblich et al (2020), rather than a short run elasticity such as Baum-Snow and Han's.

The level of floorspace expenditure is chosen to match the expenditure share on housing observed for Sydney in the ABS Household Expenditure Survey (HES)– 0.3 (1,778 Sydney observations). The ratio of expenditure share by type is calibrated to expenditure shares of disposible income by education level in Sydney by type observed in the Survey of Income and Housing (SIH, 2,212 sydney observations) 2015-16, at 0.046. This gap is then split around the mean of 0.3, to arrive at expenditure shares of 0.277 for high types and 0.323 for low types. [7]

Calibration of the remaining parameters is described in detail below. In short, the commute cost parameter ν is calibrated via the gravity equation estimated using census origin-destination commute data; epsilon is then calibrated by matching variation of wages in the data with variation of adjusted wages from commute market clearing; productivity and graduate intensity is calibrated to match firms' zero profit condition given observed wages and rents; amenities are calibrated to match the model's location shares given observed residence counts, rents, and workplace market access (from observed workplace counts and the estimated commute cost).

Table 1 lists the parameters of the static model, their values, and how each is set.

[7] The expenditure share in the HES is chosen for the level because total expenditure share matches the model object better than disposible income share.

Table 1: Static model parameters. Values for the two-type model; the one-type model shares them, with a single housing expenditure share $1-\beta = 0.30$ and commuting decay $\nu = 0.1264$ per minute. Location-specific objects are on 400 m hexagons in the Sydney SUA, recovered by inverting the model at the observed 2016 equilibrium.

	Description	Value	Source / calibration
Taken from the literature			
$1-\alpha$	Firms' floorspace cost share	0.20	Valentinyi and Herrendorf 2008, as in Ahlfeldt et al. 2015
ζ	Floorspace supply elasticity	1.86	Construction land share 0.35, Combes et al. 2021 , consistent with Sydney rents - land values elasticity
λ	Productivity spillover elasticity	0.07	Ahlfeldt et al. 2015
δ	Productivity spillover decay (per min)	0.36	Ahlfeldt et al. 2015
η	Amenity spillover elasticity	0.15	Ahlfeldt et al. 2015
ρ	Amenity spillover decay (per min)	0.76	Ahlfeldt et al. 2015
Calibrated to Sydney data			
$1-\beta^h,\ 1-\beta^l$	Housing expenditure shares, high / low	0.277 / 0.323	HES 2015–16 (level); SIH 2015–16 by education (type gap)
$\nu^h,\ \nu^l$	Commuting decay $\nu^t = \varepsilon\kappa^t$ (per min)	0.1117 / 0.1342	PPML gravity on 2016 census commuting flows
ε	Fréchet shape	6.2	Dispersion of observed vs. adjusted wages
$\kappa^h,\ \kappa^l$	Commuting cost (per min)	ν^t/ε	Implied by ν^t and ε
χ	Low- to high-type wage level	0.6073	2016 census workplace incomes by education
$L^h_{tot},\ L^l_{tot}$	Workers of each type	0.84m / 1.15m	2016 census
Location-specific			
$a^h_j,\ a^l_j$	Fundamental productivity	by location	Zero-profit condition, net of spillovers
$b^h_i,\ b^l_i$	Fundamental amenity	by location	Residential choice, net of spillovers
ψ_i	Construction productivity	by location	Floorspace market clearing
$T_i,\ E_j$	Fréchet residence / workplace scales	1	Absorbed into adjusted variables (Section 3.2)
K_i	Developable land	by location	Urban-zoned land net of water and undevelopable mesh blocks
τ_{ij}	Travel time (min)	by pair	Car routing on 2021 OpenStreetMap (r5r), motorways opened 2016–21 closed

3.2 Adjusted Variables

The Fréchet residence scale parameter T_i and workplace scale parameter E_j enter the model isomorphically with amenity and productivity for the equations that I use to calibrate the model, so I follow Ahlfeldt et al in working with the following adjusted variables, identified by a tilde:

$$\tilde{w}_j^t \equiv w_j^t E_j^{1/\varepsilon^t} \tag{13}$$

$$\tilde{A}_j^t \equiv A_j^t E_j^{\alpha/\varepsilon^t} \tag{14}$$

$$\tilde{B}_i^t \equiv B_i^t T_i^{1/\varepsilon^t} \tag{15}$$

See the Appendix for more details. Finally, I follow Ahlfeldt et al in working in variables normalised by their geometric mean. In the closed city there is no outside option anchoring the absolute level of fundamentals, so only relative levels are identified and the geometric mean is a labelling choice; all counterfactual comparisons are reported as ratios and are invariant to it.

3.3 Gravity

Commute-time responsiveness is driven by the parameter $\nu = \varepsilon\kappa$. A large κ directly increases the utility costs of a long commute. A large ε makes the Fréchet distribution more compact, and therefore an individual's idiosyncratic utility draws will be more similar from one residence-workplace location pair to another. With idiosyncratic preferences mattering less, workers are more responsive to locations' fundamental characteristics, including commute costs.

To estimate ν, I use commute data from the 2016 census and my transport network derived travel time matrix. Because the ABS censors small count units for privacy reasons, I would systematically under-count the share of workers choosing long distance commute pairs, which would cause me to overestimate the travel costs ν. To reduce this error, I use data aggregated to the larger SA2×DZN level so there are fewer small counts. The tradeoff of choosing a larger geographic unit is that the travel time matrix assesses travel time from the unit's centre, and so commutes just over the border from one unit to another will be assigned the same travel time as a trip to the far end of the neighbouring unit. Because workers prefer shorter commutes, this will lead us to systematically overestimate log travel time, and more so for closer SA2 pairs. To check the impact of any errors resulting from these measurement issues, I run robustness checks with alternate values of ν.

From the commute shares equation, we can see that the number of individuals commuting between locations has the typical gravity form — it depends on a residence effect, a workplace effect, a constant, and a function of travel time:

$$L_{ij} = \frac{L_{\text{tot}}}{\Phi} \left(T_i \left(B_i \right)^{\varepsilon} Q_i^{-(1-\beta)\varepsilon} \right) \left(E_j \left(w_j \right)^{\varepsilon} \right) e^{-\nu\tau_{ij}} \tag{16}$$

Which suggests the following regression equation for ν, with the residence fixed effect $\mathcal{R}_i$ absorbing $\{T_i, B_i, Q_i\}$ and the workplace fixed effect $\mathcal{F}_j$ absorbing

$\{E_j, w_j\}$:

$$L_{ij} = \exp(\mathcal{R}_i + \mathcal{F}_j - \nu\tau_{ij})\,\eta_{ij}, \qquad E[\eta_{ij} \mid \tau_{ij}, \mathcal{R}_i, \mathcal{F}_j] = 1 \tag{17}$$

which I estimate via PPML regression, separately for each type.

I find that $\hat{\nu} = 0.1264$ per minute (standard error 0.0019, clustered by residence and workplace), estimated on all 559,980 pairs of an SA2 of residence and a DZN of work inside the Sydney SUA, with zero flows retained. Holding fixed the characteristics of both ends of the commute, the flow between a pair halves with every 5.5 minutes of travel time. For the two worker types, commuting flows by education are only available at the coarser SA2×SA2 level, where the coarse workplace geography flattens the estimated decay. I therefore take the ratio of the two types' decay rates from type-by-type estimates on that table, and set their level so that the pooled estimate on the same table matches $\hat{\nu}$. This gives $\hat{\nu}^h = 0.1117$ for high education workers and $\hat{\nu}^l = 0.1342$ for low education workers: the commuting of low education workers falls off faster with travel time.

Figure 4 checks the exponential form. It replaces the linear travel time term in equation (17) with indicators for five-minute travel time bands, estimated by the same PPML regression on the same sample, and plots each band's coefficient against the band's mean travel time. Between 5 and 45 minutes, the range that contains 85 per cent of commuters, every band lies within 0.3 log points of the exponential. The fit is worst at the two ends. Commutes under five minutes are about 0.9 log points more common than the exponential predicts, but 86 per cent of that band's flow goes to workplaces inside the commuter's own residence SA2, so the excess reflects people working in their home area rather than the cost of the first few minutes of travel. Beyond about 35 minutes the decay flattens, and long commutes are increasingly more common than the exponential predicts; commutes of an hour or more, however, are only 1.3 per cent of flows. I keep the exponential form of the model, as in Ahlfeldt et al 2015 [3].

3.4 Adjusted wages from commute market clearing and frechet shares

Worker shares respond not to raw wages but to the 'transformed wage' $\omega_j^t \equiv (\tilde{w}_j^t)^{\varepsilon^t} = E_j (w_j^t)^{\varepsilon^t}$, ie the wage scaled by the local frechet shifter (amenity of working in location j) and the frechet scale parameter. We can solve for the value of these adjusted wages through the commute market clearing equation:

$$L_{Fj} = \sum_i \frac{\omega_j\, e^{-\nu\tau_{ij}}}{\sum_s \omega_s^t\, e^{-\nu\tau_{is}}}\, L_{Ri} \tag{18}$$

These are S equations in the S unknowns $\{\omega_j\}$, and we solve for the adjusted wages via a fixed point algorithm.

With the transformed wages in hand, we can calibrate ϵ using the same procedure as Ahlfeldt et al 2015. Under the assumption that varition in E_j is

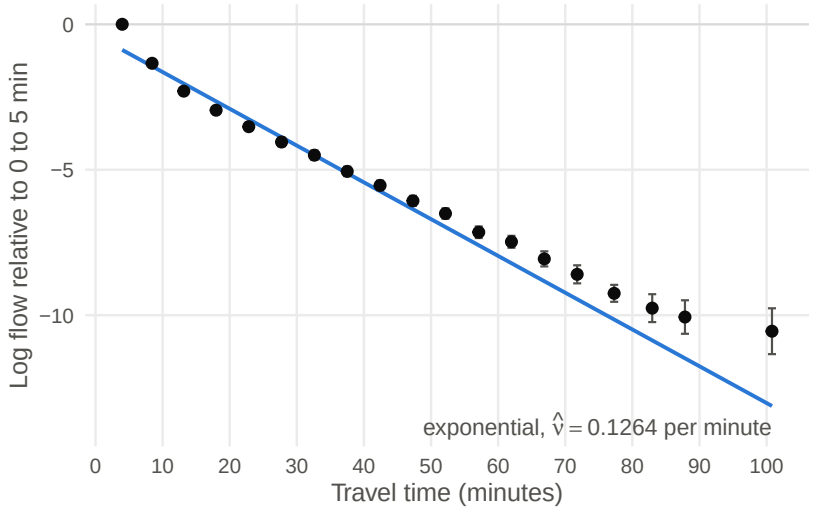


Figure 4: Commuting flows by travel time. Points are coefficients on five-minute travel time bands, relative to the 0–5 minute band, from a PPML regression of 2016 census commuting flows on the band indicators and residence and workplace fixed effects, with 95 per cent confidence intervals clustered by residence and workplace. The last point pools all pairs of 90 minutes or more. The line is the exponential decay at the estimated $\hat{\nu} = 0.1264$ per minute; the fixed effects absorb its level, so it is drawn at the flow-weighted mean of the band coefficients. Pairs of an SA2 of residence and a DZN of work inside the Sydney SUA, congested network travel times.

uncorrelated with wages and small compared to variation in wages, the value of epsilon relates variance in observed wages to variance in transformed wages:

$$\sigma^2_{\ln w_i t} = (1/\epsilon)^2 \sigma^2_{ln\omega_{it}}$$

Which gives a value of 6.2, close to the value in Ahlfeldt et al 2015 of 6.83 for Berlin.[8]

3.5 From land values to floorspace rents

The NSW Valuer General provides assessed unimproved land values for all properties in NSW. Agents in the model, however, respond to the rental price of floorspace. I follow Ahlfeldt et al in using the construction sector model to convert from land values to rental prices.

First we must convert from land values V_i (stocks) to land rental prices R_i (flows). If we assume a simple interest rate r that is constant across space, then under a perpetuity assumption the value is the discounted rent,

$$V_i = \frac{R_i}{r}, \qquad \text{i.e.} \qquad R_i = r\, V_i. \tag{19}$$

Given that we are going to normalise by the geomean, it follows that we do not need to take a stance on this interest rate beyond the assumption that it is constant across space.

To convert from land rental prices to floorspace rental prices, we turn to the construction model described earlier. With a perfectly competitive construction sector profits flow entirely to landlords, so the land rental price equals landlord profits per unit land. Substituting the floorspace FOC $f_i = \psi_i Q_i^{\zeta}$ into the landowner profit per unit land,

$$\begin{aligned} R_i &= Q_i f_i - \frac{\zeta}{1+\zeta} \psi_i^{-1/\zeta} f_i^{1+1/\zeta} \\ &= Q_i \psi_i Q_i^{\zeta} - \frac{\zeta}{1+\zeta} \psi_i^{-1/\zeta} (\psi_i Q_i^{\zeta})^{1+1/\zeta} \\ &= \psi_i Q_i^{1+\zeta} \left(1 - \frac{\zeta}{1+\zeta}\right) \\ &= \frac{1}{1+\zeta} \psi_i Q_i^{1+\zeta} \end{aligned} \tag{20}$$

[8] Because the commute market clearing equation is solved at the hexagon level but our wage data comes at DZN/SA2 level, we need to aggregate the model inferred wages to a comparable level. Because some hexagons contain multiple DZN's, the hexagon level model wages wont be able to match the variance in DZN level wages. We therefore aggregate hexagon wages up to the SA2 level to ensure multiple hexagons per aggregated unit, and hence that model implied variation at the differing sized SA2 unit levels is driven by real variation rather than aggregation measurement errors. Note that the best option within a group of workplaces is also frechet distributed with the same shape parameter, so within the model it is the same ϵ that drives commute clearing responsiveness to wages and therefore the calibrating equation is consistent at the aggregated level

Combining (19) and (20) yields

$$V_i = \frac{\psi_i Q_i^{1+\zeta}}{r(1+\zeta)} \qquad \iff \qquad Q_i = \left(\frac{r(1+\zeta)\, V_i}{\psi_i} \right)^{\frac{1}{1+\zeta}} \tag{21}$$

Now because we only observe V_i, and I do not have an accurate measure of building heights for the whole of Sydney, I am forced to follow Ahlfeldt et al 2015 in assuming a constant $\psi_i = \psi$ for the sake of this conversion. In the calibration, I later allow ψ_i to vary across cells once more, in order to ensure floorspace market clearing. As in Ahlfeldt et al 2015, this introduces a tension with our assumption of a constant ψ here. In robustness checks I plan to compare inferred floorspace to lidar building height data for central Sydney, which will serve as an over-identification check (the model cannot perfectly match observed floorspace, rental prices and floorspace market clearing, because it is underparametrised relative to the real world factors determining construction heights, so the comparison with observed building heights is a check of how accurately the model nonetheless captures floorspace construction).

The common ψ and common r cancel from every geomean-normalised object, so we solve for rental prices via:

$$\frac{Q_i}{\bar{Q}} = \left(\frac{V_i}{\bar{V}} \right)^{1/(1+\zeta)} \tag{22}$$

The census also provides a measure of rental payments by renting households. Renting households are a selected sample, and total rent is not the same object as the rent per unit of floorspace in our model, which is why I prefer my model derived conversion from more accurately observed unimproved land values per unit area. However, the rental payment data does allow for a reassuring check. Using the 2016 census, I first regress log weekly rent per renting dwelling on SA2 fixed effects and on dwelling structure × number-of-bedrooms fixed effects [9]. The SA2 fixed effects are then a rent index that compares dwellings of similar structures across the city. Regressing this index on log land value per unit area across 247 SA2s gives the cross-sectional elasticity of rents with respect to land values. Provided floorspace per dwelling of a given structure and bedroom count does not vary systematically across SA2s, the index moves one-for-one with $\ln Q_i$, and the model equation (22) implies that this elasticity should be

$$\frac{d \ln Q_i}{d \ln V_i} = \frac{1}{1+\zeta}, \tag{23}$$

which is 0.350 in the model at my headline $\zeta = 1.86$. In the data we find an elasticity of 0.340 (SE 0.024), which inverts to $\zeta = 1.94$ with a 95% confidence interval of $[1.58, 2.42]$.

[9] (structures: separate house; semi-detached or terrace house of one, or of two or more, storeys; flat in a block of one to two, three, or four or more storeys; flat attached to a house; bedrooms from none to five or more)

3.6 Graduate intensity from relative wages

Graduate intensity is inferred from what graduates earn relative to non-graduates in each location, given how many of each type work there. Rearranging the wage equation (9),

$$\mu_j^h = \frac{\tilde{w}_j^h}{\tilde{w}_j^l}\Big(\frac{L_{Fj}^h}{L_{Fj}^l}\Big)^{1/\sigma} \tag{24}$$

Intuitively, if graduates in a location earn a high wage relative to non-graduates even though many graduates already work there, firms in that location must value graduate labour highly, so graduate intensity is high. The elasticity of substitution σ sets how much weight the employment mix gets: the more easily one type can stand in for the other, the less a large graduate workforce pushes down graduates' relative wage, and the closer graduate intensity is to the relative wage alone. Both types share the same workplace shifter E_j, so adjusting wages leaves their ratio unchanged, and (24) gives the same μ_j^h whether it is computed from adjusted or actual wages. In locations where no graduates work, $\mu_j^h = 0$.

3.7 Adjusted productivity from zero profits

Adjusted productivity is inferred from firms' wage offers given rental prices.Taking the geomean of the terms in the zero profit condition, the multiplicative constants drop out and we find

$$\frac{\tilde{A}_j}{\bar{\tilde{A}}} = \Big(\frac{\tilde{w}_{Fj}}{\bar{\tilde{w}}_F}\Big)^{\alpha}\Big(\frac{Q_j}{\bar{Q}}\Big)^{1-\alpha}, \qquad \tilde{w}_{Fj} \equiv \frac{\sum_t \tilde{w}_j^t L_{Fj}^t}{L_{Fj}}, \quad \tilde{w}_j^t = (\omega_j^t)^{1/\varepsilon^t} \tag{25}$$

Intuitively, if wages in a location are high despite high rental prices, then productivity must also be high. With ω_j from Step 2 and Q_j from the land value conversion, this delivers $\tilde{A}_j$ for every location with positive employment.

I take $\alpha = 0.80$, the Valentinyi–Herrendorf (2008) sectoral factor income share for non-residential structures, in line with Ahlfeldt et al.

3.8 Adjusted amenities from residential choice

Adjusted amenities are inferred, type by type, from residential choices. Summing the unconditional commute shares π_{ij}^t over workplaces gives the residential population:

$$L_{Ri}^t = \frac{L_{\text{tot}}^t}{\Phi^t}\, T_i\, (B_i^t)^{\varepsilon^t}\, Q_i^{-\varepsilon^t(1-\beta^t)}\, W_i^t \tag{26}$$

where $W_i^t \equiv \sum_s E_s (w_s^t / d_{is}^t)^{\varepsilon^t} = \sum_s \omega_s^t\, e^{-\nu^t \tau_{is}}$ is type-t commuting market access. The term L_{tot}^t/Φ^t is a constant across locations, so differencing from the geomean removes it:

$$\frac{\tilde{B}_i^t}{\bar{\tilde{B}}^t} = \left(\frac{L_{Ri}^t}{\bar{L}_R^t}\right)^{1/\varepsilon^t} \left(\frac{Q_i}{\bar{Q}}\right)^{1-\beta^t} \left(\frac{W_i^t}{\bar{W}^t}\right)^{-1/\varepsilon^t} \tag{27}$$

Intuitively, if a large number of people of a type live in a location despite high rental prices and low access to jobs that pay that type well, it must have high amenities for that type.

3.9 Local construction suitability from floorspace market clearing

Local construction productivity is calibrated from floorspace market clearing. Total floorspace expenditure $\tilde{X}_i$ is the $(1-\beta^t)$ share of each type's resident labour income plus the $(1-\alpha)/\alpha$ share of each type's workplace wage bill which equals firms' floorspace spending under CRS.

$$\tilde{X}_i \equiv \sum_{t\in\{h,l\}} \Big[(1-\beta^t)\Big(\sum_j \pi^t_{j|i}\,\tilde{w}^t_j\Big)L^t_{Ri} + \frac{1-\alpha}{\alpha}\,\tilde{w}^t_i\,L^t_{Fi} + (1-\beta^t)\,\tilde{w}^t_P\,L^t_{PRi} + \frac{1-\alpha}{\alpha}\,\tilde{w}^t_P\,L^t_{PFi}\Big], \qquad \pi^t_{j|i} = \frac{\omega^t_j\, e^-}{W} \tag{28}$$

Equating floorspace income and expenditure,

$$\tilde{\psi}_i K_i Q_i^{1+\zeta} = \tilde{X}_i \tag{29}$$

and rearranging in geomean-normalised form, we have the calibration equation:

$$\frac{\tilde{\psi}_i}{\bar{\tilde{\psi}}} = \left(\frac{\tilde{X}_i}{\bar{\tilde{X}}}\right)\left(\frac{K_i}{\bar{K}}\right)^{-1}\left(\frac{Q_i}{\bar{Q}}\right)^{-(1+\zeta)} \tag{30}$$

3.10 Amenity and productivity fundamentals from agglomeration

Currently I do not have a means of estimating the agglomeration spillover parameters. Instead, I take the values from Ahlfeldt et al at baseline, and test the sensitivity of results to other values in robustness checks. Taking the agglomeration parameters as given and the adjusted productivity and amenity values from the calibration above (with L_{Fs}, L_{Rs} the observed baseline totals), we solve for the corresponding adjusted fundamentals:

$$\tilde{a}_j = \tilde{A}_j \left(\sum_s e^{-\delta\tau_{js}}\frac{L_{Fs}}{K_s}\right)^{-\lambda} \tag{31}$$

$$\tilde{b}_i = \tilde{B}_i \left(\sum_s e^{-\rho\tau_{is}}\frac{L_{Rs}}{K_s}\right)^{-\eta} \tag{32}$$

4 Benchmarking the costs of the lopsided employment distribution

With the model parameters calibrated to the baseline equilibrium, I can alter features of the model, solve for counterfactual equilibria, and consider the impacts on output, welfare, commute times and other outcomes for the city. In

this section, I run a series of stylized counterfactuals to try to understand the costs of Sydney's lopsided jobs distribution. In an appendix, I report several other counterfactuals aimed at understanding other aspects of Sydney's shape: the cost of the harbor (as a development and transport barrier) and of its non-circular urban boundaries.

4.1 Shifting the central business district

To give an indication of the cost of Sydney's lopsided jobs distribution, I compare outcomes under the counterfactual equilibrium in which the fundamentals of the current CBD are shifted to a more central location. Operationally, I take the fundamental productivity, amenity, construction productivity and build-able area from the 9 square kilometers closest to the center of Sydney's cbd, and swap them with the corresponding fundamentals around Parramatta, holding the transport network constant (See Figure 5 for maps of the switched fundamentals).

The city's total area, average fundamentals, and transport network are unchanged by this swap, so the counterfactual captures just the impact of changing where in the city the CBD's high productivity land is located. This counterfactual is stylized: part of the reason for the high productivity in the CBD (along with historical investments and coordination effects) is its harbor-side views and cool temperatures, which are not available in Paramatta. The idea is not to measure a specific real world possibility, but to isolate the model inferred impacts of the distribution of jobs being so lopsided.

In response to the shift in fundamentals, firms change their wage offers, workers change their location decisions, and developers change their construction and rent price offers. Because high productivity areas are now more accessible from more parts of the city, one would expect employment around this new center to be higher than it was around the old center. Average commute times to the center from all locations in the urban footprint fall mechanically, but the average selected commute times in equilibrium could in theory fall or rise: the mechanical effect means that for a given place of residence, the new cbd is now closer, but because the new CBD is now within acceptable travel time of more locations than the old, more people may be lured to travel to the high wage cbd instead of working locally.

4.1.1 Results: One-type

When the CBD fundamentals are shifted in the baseline specification of the basic model, output increases by 11.0%, commutes fall by 2.1% (from 20.97 to 20.54 minutes), and worker welfare increases by 2.5%. In this as in all counterfactuals, total rent take increases by the same amount as output: Under CRS + zero profits commercial rents are $(1-\alpha)Y$ and residential rents $(1-\beta)*\alpha*Y$, so the rent bill is $(1-\alpha)+\alpha(1-\beta) = 0.44$ of output in every equilibrium. Employment is also more concentrated in the counterfactual city: the number of people in the highest employment 400m hexagon increases by 50% (from 98,000 to 147,000),

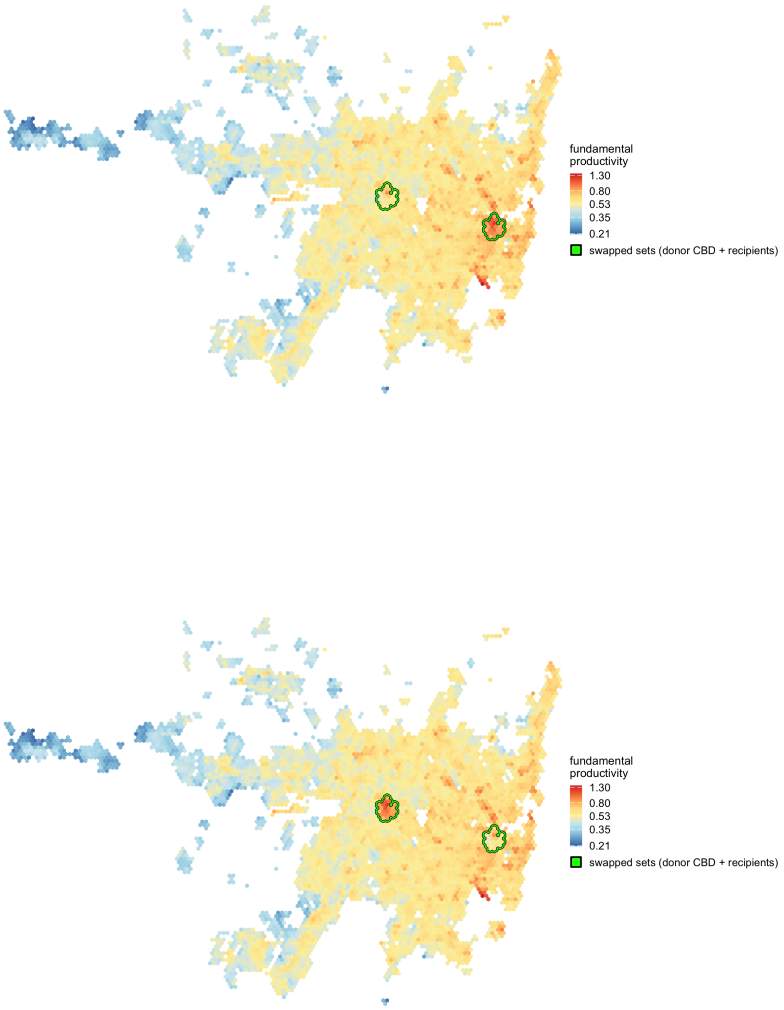


Figure 5: Fundamental productivity $\tilde{a}$ by 400 m hexagon in the baseline calibration (top) and the CBD-shift counterfactual (bottom), Sydney SUA, 2016. The two panels share a single log colour scale. The green outlines mark the two swapped 22-cell sets — the donor CBD in the city's east and the recipient cells around Parramatta Square; the counterfactual swaps the two sets' fundamentals (productivity shown here; amenity, construction productivity and developable area move with it), leaving the rest of the city and the transport network unchanged.

and as shown in Figure 8, the employment share of the highest n employment hexagons is higher in the new counterfactual for all n — there is no crossing point.

Overall, these results suggest that the more centrally located CBD leads to better commute accessibility, reducing the drag of commute costs on welfare and driving output up both directly by increasing employment in more productive areas and indirectly via the increased productivity spillovers resulting from higher density in these areas. In other words, sorting of CBD workers and agressive densification around the center do not appear to be able to fully counteract the geographic isolation faced by the CBD's eastern harborside location.

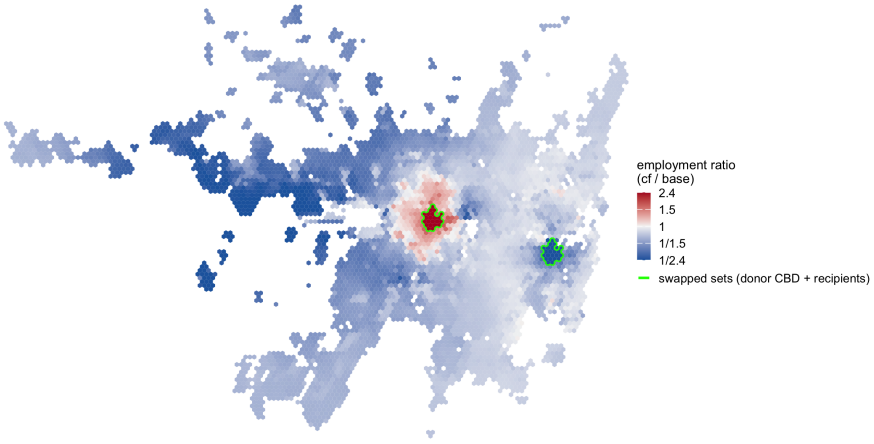


Figure 6: Change in workplace employment under the CBD shift: the ratio of counterfactual to baseline employment by 400 m hexagon (log scale; extremes clipped at the ends of the colour range). The green outlines ring the two swapped cell sets. Employment re-forms around Parramatta — the recipient cluster rises from 44,000 to 736,000 workers, with a halo of gains in the surrounding cells — while the vacated donor CBD collapses from 403,000 to 17,000 workers and drags its neighbourhood down with it. Outside the two swap sets employment drains from almost the whole city (the median cell loses 25% of its workers and only 4.5% of cells gain), most visibly in the far west.

To see which of the model's responses deliver this gain, Table 2 builds the counterfactual up one response at a time, starting from the smallest response the city can make and ending at the headline. In the first row nobody changes

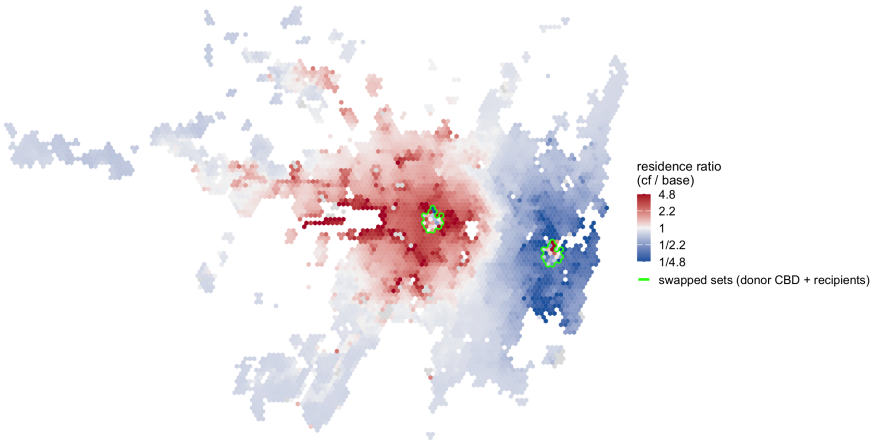


Figure 7: Change in residence under the CBD shift: the ratio of counterfactual to baseline residents by 400 m hexagon (log scale; extremes clipped at the ends of the colour range). The green outlines ring the two swapped cell sets. Residence re-polarises around the relocated CBD: a broad halo of gains — up to almost fivefold — forms around Parramatta through the middle of the city, while the east around the vacated CBD drains. In aggregate the Parramatta recipients gain residents (from 23,000 to 30,000) and the donor CBD loses them (from 53,000 to 45,000), with individual cells inside each set moving in both directions.

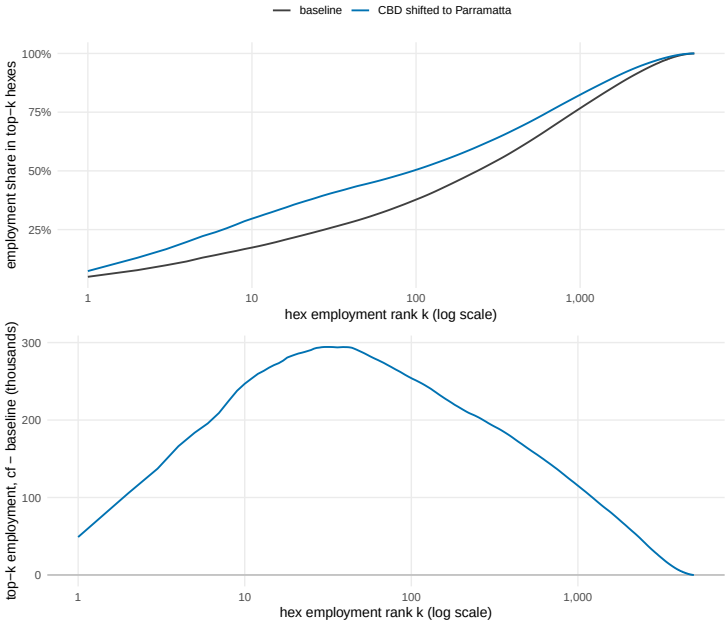


Figure 8: Employment concentration, baseline vs the CBD-shift counterfactual. Top panel: cumulative employment share of each city's top-k hexagons (each scenario ranked by its own hexagon employment). Bottom panel: the difference between the two curves. The counterfactual curve lies above the baseline at every rank — the shifted city is more employment-concentrated throughout the distribution, with the surplus peaking at 294,000 workers in the top $k = 31$ hexes (i.e. 40.6% of all workers vs 26.0% of all workers at baseline, in 12.9 km^2).

home or job and no prices or wages change, but workers and residents move with their jobs: the only change is that cbd workers location of work has moved to parramatta and vice versa (and the same for CBD residents). On this direct effect alone the shift lowers welfare by 3.7%, as the average commute rises from 21.0 to 23.5 minutes, reflecting the fact that in the old real world equilirbium CBD workers sort to live near the harborside cbd location. In the next line we move the stock of buildings and total productivity and amenity of the cbd but now let people choose where to live and work, while still holding the total building stock constant and not allowing productivity or amenity spillovers. As a result, employment centralises around the new CBD and output rises by 3.5%. However, the shift leaves welfare roughly unchanged (-0.2%) because there is no building supply response, and hence rents rise substantially because workers are trying to centralise more around the same set of cbd buildings now stuck in a surrounding area that is much less built up. Allowing the building market to respond is the largest single step: it adds 1.3 percentage points to welfare, more than half of the headline gain, so that without relying on agglomeration spillovers at all the shift raises welfare by 1.2%. Productivity spillovers then add little to welfare (0.09 percentage points) but over two percentage points to output, and amenity spillovers add the remaining 1.2 points: 0.35 at roughly half the elasticity of Ahlfeldt et al (2015) [3] (their reunification-only estimate, $\eta = 0.076$) and a further 0.89 at the full value. The amenity step is convex, consistent with the residence-amenity feedback being self-reinforcing under elastic floorspace at the full elasticity ($\eta(1+\zeta) = 0.43 > 1-\beta = 0.30$) but not at half strength (0.22).

Though the direction of the results are robust to alternate parameter values, the magnitudes are somewhat sensitive:

4.2 Inequality

To understand the effects on inequality, we turn to the model with workers split by type. In the counterfactual equilirbium, worker-renter welfare for high types increases 3.7% for graduates but falls by 0.1% for non graduates (note that this does not include landlord profits/forces homeowners to pay for housing they may own, so if we redistributed landlord profits the overall impact would be positive for both but still favor graduates). The reason is that with the CBD in a more accessible location, employment centralises more around the CBD than it did in its old location. As the CBD has the most graduate intensive production technology, this centralisation favours high types.

This high-type favourability happens despite the new CBD's location improving CBD access more for low types (as seen for example in the first line of the results breakdown, which shows that before workers are allowed to resort, the welfare loss for low types is smaller than that for high types, because more low types made long journeys from the west to the old cbd location than high types did). The model explains the substantial sorting of graduates east and non graduates west (the median non graduate is 17km West of the CBD, vs 9km for gradutes) primarily as a result of the greater relative amenity for high

Table 2: CBD shift built up one response at a time (Sydney SUA, 400 m hexagons, 2016 calibration on 2016 roads; $\varepsilon = 6.2$, $\zeta = 1.86$). Every row moves the same bundle of fundamentals and adds one response to the row above; the last row is the headline. Changes are relative to the baseline equilibrium; the increment is the change in welfare over the row above, in percentage points.

	Δ welfare	Increment	Δ output	Δ commute
Direct effect: nobody changes home or job	−3.67%		0.00%	+12.02%
+ people choose where to live and work	−0.17%	+3.51	+3.50%	−0.90%
+ building supply responds	+1.16%	+1.33	+4.88%	−0.68%
+ productivity spillovers	+1.25%	+0.09	+7.06%	−0.05%
+ amenity spillovers at half strength ($\eta = 0.076$)	+1.60%	+0.35	+8.19%	−0.35%
+ amenity spillovers at full strength ($\eta = 0.15$): headline	+2.49%	+0.89	+10.99%	−2.06%

Notes: In the first three rows each location's total productivity and amenity are held at their baseline levels and move with the swapped fundamentals (no agglomeration spillovers); in the first two rows each location's stock of floorspace is also held fixed and moves with the swapped locations. In the first row every worker and resident moves with their building, so wages and rents do not change, and welfare holds each worker's choice of home and workplace fixed. The increments depend on the order in which responses are added: switched on after amenity spillovers rather than before, productivity spillovers add 0.33 percentage points rather than 0.09, and amenity spillovers 1.01 rather than 1.24. Holding residence amenities fixed in a different way — letting fundamental amenity move with the swapped locations while each location keeps its baseline amenity spillover — gives a welfare gain of 1.29% (output 6.99%), against 1.25% in the fourth row.

Table 3: CBD shift: robustness of the headline, with all of the model's responses switched on (Sydney SUA, 400 m hexagons, 2016 calibration on 2016 roads; $\zeta = 1.86$ and $\varepsilon = 6.2$ unless varied). The specification brackets change which of the CBD's fundamentals move with it; each parameter variation changes one parameter and recalibrates all location fundamentals before re-solving. The amenity spillover elasticity $\eta = 0.076$ appears as a row of Table 2. The $\varepsilon = 8$ variation is omitted: its baseline calibration fails the round-trip acceptance gate, so its counterfactuals are not valid.

	Δ welfare	Δ output
Headline	+2.49%	+10.99%
Specification brackets		
Developable area left in place (keepK)	+2.41%	+10.52%
Area and construction productivity in place (keepKψ)	+1.60%	+3.45%
Parameter variations (one at a time)		
Amenity spillover decay $\rho = 0.38$	+2.51%	+10.62%
Floorspace supply elasticity $\zeta = 0.5$	+0.69%	+5.93%
Floorspace supply elasticity $\zeta = 3.0$	+3.98%	+17.13%
Fréchet shape $\varepsilon = 4.5$	+2.64%	+13.47%
Own-cell travel time undiscounted	+2.46%	+11.07%

types in the East.

To see what the model attributes this sorting to, take the ratio of the two types' residents in each location. Summing the commuting shares (2) over workplaces, anything common to both types cancels, and since the amenity spillover in (E9) depends on total residents it cancels too, so that $B_i^h/B_i^l = b_i^h/b_i^l$. This leaves an exact identity, which holds cell by cell in the calibrated model:

$$\ln \frac{L_{Ri}^h}{L_{Ri}^l} = \text{const} + \underbrace{\varepsilon \ln \frac{B_i^h}{B_i^l}}_{\text{amenity}} + \underbrace{\varepsilon(\beta^h - \beta^l) \ln Q_i}_{\text{housing share} \times \text{rent}} + \underbrace{\ln \frac{W_i^h}{W_i^l}}_{\text{access to own-type jobs}}, \qquad W_i^t = \sum_j \Big(\frac{w_j^t}{d_{ij}^t} \Big)^{\varepsilon}. \tag{33}$$

An amenity both types value equally, such as the harbour, sorts no one: under Cobb–Douglas preferences its rent premium costs both types the same share of income. Income levels cancel too, so only where each type's jobs are, relative to the other type's, enters the access term. Production does not enter at all, so the decomposition is the same for any σ. Figure 9 maps the three terms. Averaging over residents, the log ratio of graduates to non-graduates is 0.96 higher east of Parramatta than west of it. The amenity term accounts for 1.00 of this, the housing-share gap times rent for 0.12, and access to own-type jobs for -0.15. Access works against the pattern: graduates' flatter commuting decay ($\nu^h < \nu^l$) makes distance from jobs cheaper for them, which favours them on the far western fringe. The east–west amenity difference is equivalent to graduates being willing to give up 17% of their income, relative to non-graduates, to live east of Parramatta.

Table 4: CBD shift by worker type, built up one response at a time (two-type model with CES production, $\sigma = 1.4$; Sydney SUA, 400 m hexagons, 2016 calibration on 2016 roads; $\varepsilon = 6.2$, $\zeta = 1.86$, $1 - \beta^h = 0.277$, $1 - \beta^l = 0.323$). Every row moves the same bundle of fundamentals and adds one response to the row above; the last row is the headline. Welfare is the expected utility of all workers of each type; commutes are each type's average, against its own baseline.

	Δ welfare		Δ output	Δ commute	
	Grad.	Non-grad.		Grad.	Non-grad.
Direct effect: nobody changes home or job	−5.07%	−2.47%	0.00%	+17.77%	+7.93%
+ people choose where to live and work	+1.27%	−2.04%	+2.34%	+0.13%	−1.74%
+ building supply responds	+2.17%	−0.79%	+2.94%	+0.26%	−1.36%
+ productivity spillovers	+2.22%	−0.58%	+3.93%	+0.79%	−1.14%
+ amenity spillovers at half strength ($\eta = 0.076$)	+2.56%	−0.40%	+4.68%	+0.40%	−1.11%
+ amenity spillovers at full strength ($\eta = 0.15$): headline	+3.73%	−0.09%	+6.44%	−1.79%	−1.61%

Notes: Rows are defined as in Table 2. Public-sector workers are ordinary workers who choose where to work. The non-graduate loss appears in the second row, before any supply or spillover response: once workers re-sort, the relocated CBD's graduate-intensive technology bids up the graduate premium, and every later response returns part of that loss but not all of it. With perfectly substitutable types ($\sigma = \infty$) the same rows give −5.07/−2.47, −0.55/−0.07, +0.45/+1.39, +0.60/+1.53, +0.78/+1.92 and +1.55/+2.81 per cent (graduates/non-graduates). Switched on after amenity spillovers rather than before, productivity spillovers add +0.44 and +0.11 percentage points.

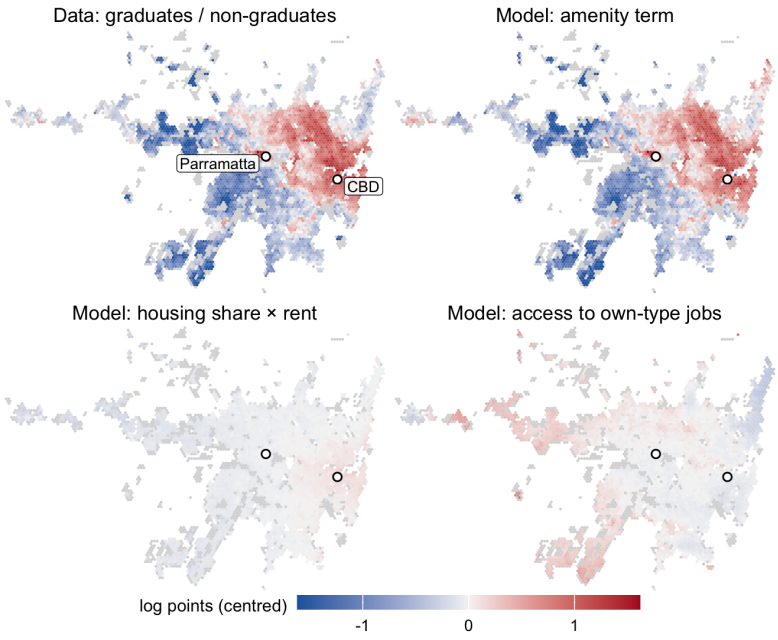


Figure 9: Why graduates and non-graduates live apart. Top left: the log ratio of graduate to non-graduate residents in each 400 m hexagon, 2016 census. The other three panels are the three terms of equation (33) in the calibrated two-type model, which add up to the data. Each panel is centred on its residents-weighted mean, on one shared colour scale (red: more graduates than the city average). Grey: one type absent from the cell (3,786 of 4,753 cells have both).

5 Policy analysis: Government Office Moves

The above counterfactuals suggest that the costs of Sydneys distorted urban form may be substantial and contribute to inequality. How can policymakers respond? One approach is to counter the city's lopsided employment distribution by moving the public service jobs that they directly control. In 2013, the NSW government launched a policy to do just this: the 'decade of decentralization' shifting NSW public service jobs out of the CBD with Parramatta (20km WNW of the CBD) emerging as the main location of new offices. To estimate the effects of this policy in reduced form, I run a Synthetic Control analysis of the impacts on land values, private office construction approvals, and employment.

5.1 Background

The "Decade of Decentralisation" was a 2011 Coalition election committment to move government offices from the CBD and into greater Sydney and regional NSW. Parramatta became the central hub for new government offices: From 2013- 2024, between 11,000 and 15,000 public service jobs occupying over 100,000 sqm of office space were newly housed in Paramatta, against a baseline total public and private employment of around 36,000 in 2011. Parramatta is located 20km WNW of the CBD, close to the mean place of residence of Sydney Urban area residents, and was already the second largest hub for office space and civil servants, though as stated above the new civil servants represented a major increase in local employment. This policy was partly aimed at simply reducing exposure to expensive government rents and placing civil servants closer to the people they serve. But statements in planning documents and press releases suggest it was at least partly motivated by the desire to make jobs easier to access both for civil servants themselves and for other Western Sydney residents, by encouraging private buisinesses to locate around the new government offices and reinforce Parramattas position as a 'second CBD'.

Timeline of the decentralization policy:

- Earlier moves: 2003 (Police HQ 1,300 jobs), 2009: Sydney Water HQ (1,000 jobs) and Parramatta Justice Precinct
- 2011 election commitment to 'decade of decentralisation' followed by 2013 report of 'decentralization taskforce' and government planning for office moves to Parramatta in response.
- Q4 2016: Biggest Lease signed in December(4 PSq, $> 60,000$sqm), with building beggining in 2017; Dept of education move announced in february.
- 2018: Opening of Department of education HQ (1,800 jobs (though not all moved until 2020))
- Late 2019-2020: Numerous departments to 4Psq with 5,000 government jobs

- 2022: Again, numerous departments to another purpose built high rise at 6 PSQ tower with 5,000 government jobs

5.1.1 Confounders

A key difficulty for analyzing the impact of the policy is the impact of Covid. From 2020 to late 2021, Sydney had rolling lockdowns with heterogeneous geographic incidence. This coincides with the opening of the largest new government offices at 4 Parramatta Square (occupied November 2019 - June 2020, c. 5,000 staff). Furthermore, an additional c. 4,000-5,000 staff moved into tower at 6 Parramatta Square after the census in 2022. The census was taken on 10 August 2021, when Paramatta had especially harsh lockdowns. The ABS attempted to minimise the impact on results by giving the instruction[10]: "For people who are working from home temporarily due to COVID-19, write the workplace address they would normally be at." The ABS's own validation[11] found the share of respondents whose place of work equalled their home location rose only from 7.3 percent (2016) to 8.4 percent (2021), i.e. most respondents followed the instruction – but the ABS also cautions that "people may have put their home address as the workplace address despite instructions", and place-of-work non-response rose from 10.0 percent to 11.1 percent. Perhaps more importantly, many firms that might otherwise have opened in response to the presence of new civil servants may not yet have had time to do so before covid hit. While comparing outcomes with other parts of Sydney also subject to covid through my synthetic control method may partly help to control for the covid effect, we might expect that the signal we are looking for– new firms locating in Parramatta– is likely to be instantly damped by covid, whereas any negative effect on existing firms in other areas would be damped or delayed by furlough payments and sunk costs. Additionally, the government job moves coincided with a number of other place based policies: a light rail network connecting Paramatta to the West-mead health district and surrounding neighborhood was constructed from 2020-2024. In November 2016 a major metro extension/upgrade was also announced but is not expected to open until 2032. My synthetic control analysis should therefore strictly be understood as analysing the combined effect of all these policies, though I believe a $> 20\%$ increase in local employment to be the most important component, I will attempt to net out the effect of the light rail transport in my model, and the control units are chosen from a set of government designated strategic centers that also received positive local policy shocks such as transport in this period but not the large increase in government jobs.

5.2 Data

I estimate the impact of the policy on three main outcomes: land values, office construction approvals, and employment. To measure land values I use valua-

[10](ABS Census Dictionary 2021, variable POWP)
[11]("Understanding Place of work data", information paper, 12 Oct 2022)

tions provided by the NSW valuer General. The data includes yearly valuations of all properties from 2012 -2025. They are estimates of the unimproved land value: assessors' estimates of what the land would sell for if it were empty of any buildings and available for the highest value use subject to local zoning restrictions (though they are also intended to account for rezoning potential when it is deemed likely in the near future). The assessments are organized by the valuer general for the purpose of land and council taxes, and property owners can object to a valuation if they think it innaccurate. [12]

My units of analysis for the land value series are 800m circles around Parramatta and comparable employment centers (described in the next section), so to create my time series I sum the value of properties within these circles. To understand the differing impacts on commercial vs residential markets, I categorize properties using zoning codes from the valuations data.[13] I produce four seperate indices for properties falling under different zoning groups:

1. Genuine Commercial: Only commercial centre zoning (B3/B8 pre-reform 2012-2020; E2/SP5 post-reform(2021-2025)) E2 'commercial center' and SP5 Metropolitan Centre (and corresponding earlier codes)
2. Mixed Commercial: Mixed used and other commercial codes, including the commercial center zoning above
3. Residential
4. All urban

For each of these, I take only the fixed set of properties that are present for the specific zone type definition for all of 2012-2025. This means that I drop any properties that are amalgamated, split or change zoning code, including the specific properties occupied by government offices (which were newly defined lots) and therefore the measure does not include eg properties mechanically revalued by change in zone code or by the presence of new captive government tenants. For each year, i then sum the values of just these properties. For Parramatta this leaves: all urban 725 vs the 869 in-scope properties in 2012 under the dynamic definition (i.e. 17percent of the 2012 stock excluded, dominated by the

[12] Assessors value by components: properties are grouped into a set of properties with similar size, location, and use, and given relative valuations at one point in time. The number of properties in the component vary by kind, but in general will feature a few hundered residential properties, or a few dozen small commercial properties, or individual highly valuable large commercial properties. For re-evaluations, the benchmark property is re-assessed and the other component properties values are updated by the same factor. In robustness specifications, use the universe of all property sales in greater Sydney from 2001-present, also provided by the valuer general. However, as I lack sufficient property characteristics, and the set of properties transacting changes each year, I am unable to define a wholly satisfactory local price index with this data (especially for commercial properties) and therefore primarily rely on the valuations data for my main specifications.

[13] The valuations data has property and year specific zone codes. Note the earliest available codes are from 2017, so the zone code is forward looking for pre 2017 data. For sales data, the code is taken from the nearest LV data property match, and if there is no match taken from the zoning code in the sales data. The match from pre 2012 sales to a 2017 LV zone code is c. 95 percent.

4 PSQ-era consolidations); mixed commercial 497 (vs 640); residential 226 (vs 228); genuine commercial 54 (vs 91). I produce per area specifications as raw AUD/m2 LEVELS. Currently, these are constructed by dividing the above land value sums by the area of the corresponding zones in the 2026 zoning layer provided by NSW spatial services (I am searching for earlier zoning maps). To avoid areas with rezoning shocks, I remove circle units when at least 5 percent of its in-circle, in-scope land value (over properties present at both vintages) moved across an LV vintage boundary.

As a second outcome i consider office construction approvals. The ABS provides monthly data at ABS SA2 level (around 50,000 workers, 8.5 sqkm for parramatta) of the total estimated cost of all approved office construction from 2011-2026. I subtract from this series the new office construction that was destined to be occupied by the new government jobs, leaving me with a private office construction value time series. The area and baseline number of workers varies by SA2 unit, so as the outcome variable I calculate the total value of approvals in each area per 2016 office sector employed workers, with a per area specification as robustness.

For employment, I use census data [14] from 2001, 2006, 2011, 2016, 2021, split by 1-digit ANZSIC 2006 industry codes.[15] The data is at varying geographies, which I aggregate to my circle units of analysis. In the appendix I run robustness checks of alternate area weighting designs for aggregation from census data to the circle units.

5.3 Synthetic Control Design

For each index I run a Synthetic control procedure closely following Abadie et al 2010 [1]. The aim is to create a weighted sum of comparable donor candidates that mirrors the treated unit before treatment, in the hope that they match on the underlying variables (except treatment) that drive variation in the outcome over time. The estimated treatment effect is the difference in post treatment outcomes between the treated unit and the synthetic control.

My baseline assumption is that the treatment begins in Q4 2016, when the first leases for major new office buildings in Parramatta were signed. Parramatta had already been identified as a potential key location for new office buildings in early 2013 in the first planning documents for the decentralisation policy, and the decentralisation policy itself was first foreshadowed in 2011 election commitments, so this assumption requires any land value adjustments were not significantly capitalised until lease commitments were made. My land values data currently only goes back to 2012 (and office approvals to 2011) but I am trying to access earlier land value data to rule out earlier land value responses.

[14] for 2001 and 2006 employment data is census derived but provided by transport for NSW

[15] The 2001 data (TfNSW JTW) uses ANZSIC 1993 and is crosswalked to 2006 divisions at the 4-digit level

5.3.1 Sample Construction

Synthetic control methods work best with a small, well matched donor pool. Before treatment, Parramatta was a dense employment center with a relatively large office share of employment, commercial center and genuine commercial zoning, and high rise residential buildings, relatively far from the central business district. The aim is therefore to create a pool of donors that are similar on these characteristics. To create candidates for the donor pool I first identify the high emplopyment centers by identifying all distinct 800m radius circles containing over 5,000 employees.[16] I then screen centers that have: i) a location within 3 km of the CBD (Martin Place) or of Parramatta;[17] ii) a single dominant function, with health and education (ANZSIC divisions P and Q) or industry (A–C, E and I) making up at least half of 2016 jobs; iii) a public-employment shock of their own, where the 2016–21 change in government jobs as a share of the centre's 2016 jobs, net of the same rate for Greater Sydney, is at least a fifth of Parramatta's; iv) a new Metro station (excluding conversions of existing lines) opened within 1.2 km between 2017 and 2025; or v) too few land-value observations, meaning fewer than 30 parcels of the outcome's land in some pre-treatment year (2012–16). Screens iii) and iv) remove centres that received a large place-based shock of their own after 2016, whose post-treatment paths would not show how Parramatta would have evolved without the relocation. From the 35 centres that remain, the donor pool is the 15 most similar to Parramatta: the equally weighted Euclidean distance over eight standardised pre-treatment characteristics, namely log 2016 jobs; the office-occupation, retail and hospitality (G+H) and industrial shares of 2016 jobs; the density of dwellings in high-rise flats; the commercial-core (B3/B8) share of land; the B5–B7 share of business-zoned land; and log distance to the CBD.

5.3.2 Synthetic control procedure

We have one treated unit and J donors observed over T periods, with treatment after period T_0. Let $\mathbf{Y}_1^{\text{pre}} \in \mathbb{R}^{T_0}$ be Parramatta's pre-treatment outcomes, $\bar{Y}_1$ their mean, and $\mathbf{Y}_0^{\text{pre}} \in \mathbb{R}^{T_0 \times J}$ the donors' pre-treatment outcomes. For weights $\mathbf{w} = (w_1, \ldots, w_J)'$ in the simplex $\mathcal{W} = \{\mathbf{w} : w_j \geq 0,\ \sum_j w_j = 1\}$, measure the pre-treatment fit by the root mean squared prediction error as a share of Parramatta's own level,

$$R(\mathbf{w}) \;=\; \frac{1}{\bar{Y}_1}\Big(\frac{1}{T_0}\big\|\mathbf{Y}_1^{\text{pre}} - \mathbf{Y}_0^{\text{pre}}\mathbf{w}\big\|^2\Big)^{1/2},$$

and let d_j be the distance between Parramatta and donor j in pre-treatment characteristics: the squared Euclidean distance over the standardised charac-

[16] Circles are centred on local peaks of a 50 m surface of 2016 place-of-work jobs, picked in descending order of jobs so that no two circles overlap, within the Sydney urban area.

[17] Eight established office precincts are exempt from the CBD and single-function screens. This re-admits the CBD fringe (Pyrmont–Ultimo) and Green Square–Mascot.

teristics of the donor ranking except log jobs,[18] plus the squared difference in the standardised log 2016 outcome.

The weights are chosen in two steps. First, find the best achievable pre-treatment fit, $R^* = \min_{\mathbf{w}\in\mathcal{W}} R(\mathbf{w})$, which is the standard synthetic control of Abadie et al. 2010 [1] matching on outcomes alone. Second, among all weights that fit nearly as well, choose the synthetic whose donors are closest to Parramatta:

$$\hat{\mathbf{w}} \;=\; \arg\min_{\mathbf{w}\in\mathcal{W}} \sum_{j=1}^{J} w_j d_j \qquad \text{s.t.} \quad R(\mathbf{w}) \le R^* + \delta,$$

with $\delta = 0.01$. That is, we give up at most one percentage point of pre-treatment fit, relative to Parramatta's mean level, in exchange for donors that more closely resemble Parramatta. This is the penalised synthetic control of Abadie and L'Hour 2021 [2], which minimises $R(\mathbf{w})^2 + \lambda \sum_j w_j d_j$, with the penalty λ set implicitly: the constrained solution is the penalised solution at the value of λ for which the fit constraint just binds. Fixing the fit cost δ rather than λ makes the trade-off interpretable, and comparable across units of different scale. Every placebo uses the same δ, with its own R^* and distances measured from the placebo unit. Setting $\delta = 0$ recovers the best fit alone, with any ties broken by distance; we report it as a robustness check.

The reason not to match on outcomes alone is that with 15 donors and five pre-treatment years, many quite different combinations of donors fit Parramatta's pre-treatment path almost equally well. Choosing the single best of them selects on transitory noise in a short pre-period, and can produce a synthetic that averages dissimilar centres whose differences happen to cancel. The distance penalty resolves this near-indeterminacy in favour of donors that resemble Parramatta in the characteristics that drive land values, which also limits the interpolation bias that Abadie and L'Hour 2021 [2] show arises from averaging dissimilar units.

Once we have the synthetic control weights, we estimate the treatment effect as the difference in post treatment outcomes between the treated unit and the synthetic control

Treatment effect. The estimated gap in each post-treatment year:

$$\hat{\tau}_t \;=\; Y_{1t} \;-\; \sum_{j=1}^{J} \hat{w}_j \, Y_{jt} \qquad \text{for } t > T_0$$

The identification assumption is that absent treatment the treated unit and the synthetic control would have followed the same post treatment outcome trajectory. This requires that there are no omitted variables that differentially impact treated and control units.

[18] The office, G+H and industrial shares of 2016 jobs, the density of high-rise dwellings, the commercial-core share of land, the B5–B7 share of business-zoned land and log distance to the CBD.

5.3.3 Inference

With only one treated unit, inference is difficult. Given the small number of pre-treatment periods with outcome data available, I cannot use training and testing split method of Abadie et al 2015. Instead, I use a placebo inference method following Abadie et al 2010. I remove the treated unit from the sample, run the synthetic control procedure for each of the control units in turn, and then compare the estimated placebo 'treatment' effect to the treated units estimated treatment effect. The reported p-value is the rank of the treated unit's post treatment root mean squared prediciton error among the placebo unit's post treatment root mean squared prediciton error– ie it compares how large the average post period gap is between the treated unit and its synthetic control compared to the post period gap of the control units and their placebo synthetic controls. The idea is that if the post average post treatment gap is larger for the treated than the control units, this indicates the estimation less likely to be simply due to noise/failure of the synthetic control method to provide accurate controls in this setting.

5.4 Results

5.4.1 Land values

Figure 10 shows how the donor pool is built. The candidates are every centre with at least 5,000 jobs within 800 m in 2016, outside the CBD and a 3 km buffer around Parramatta, plus the CBD-fringe office precinct. After the screens described above, the donor pool is the 15 centres most similar to Parramatta. Four of them carry the synthetic control's weight: Hurstville (0.54), Rhodes (0.28), Green Square–Mascot (0.10) and Macquarie Park (0.08); the other eleven receive none. The headline outcome is land value per m^2 of commercially zoned land (business zones, excluding industrial), measured on the same parcels throughout.

Relative to its synthetic control, commercial land value in the Parramatta circle pulls away after the December 2016 lease commitment (Figure 11). The two series track each other closely over 2012–16, with a pre-treatment root mean squared prediction error of \$44 per m^2. By 2025 Parramatta's commercial land is valued at \$4,972 per m^2 against \$3,272 for its synthetic control, a gap of +\$1,700 per m^2, or 52% of the synthetic level. Most of the gap opens between the 2019 and 2022 valuations. The 2020 and 2021 values are drawn dotted because the Valuer General set them by rule rather than from sales: a flat COVID-19 reduction of 15% to CBD office land at 1 July 2020 wherever comparable sales were lacking, held at 1 July 2021. Neither the fit nor the 2025 gap uses them. Dropping each weighted donor in turn keeps the 2025 gap between +\$1,516 and +\$1,829 per m^2, so no single centre drives it.

Figure 12 repeats the exercise with each of the 15 donors treated as if it had received the offices, each with its own synthetic control and the same fit tolerance. Thirteen of the placebos fit their own pre-period within five times Parramatta's pre-treatment error, and none of them has a larger post-treatment

gap than Parramatta, so $p = 1/14 = 0.07$, the smallest value this design can produce (none of the 15 unfiltered placebos is larger either). With 14 units the p-value is coarse, and it rests on pool membership. Under a slightly different public-jobs screen (a quarter of Parramatta's raw change rather than a fifth of its change net of Greater Sydney), Neutral Bay enters as the fifteenth donor. Its placebo gap is larger, and p doubles to $2/14 = 0.14$, while the gap itself is essentially unchanged at +$1,692.

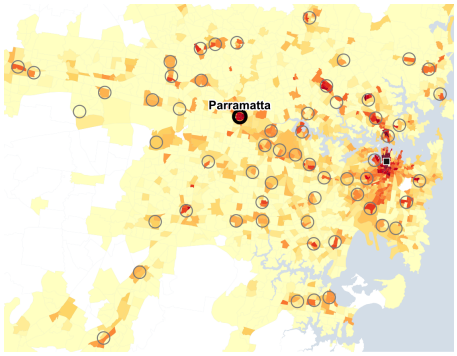
Figure 13 collects variants of the headline, each as the 2025 gap in per cent of the synthetic level. The gap is robust to the weighting rule and to the zoning definition. Choosing the best pre-treatment fit alone ($\delta = 0$) gives +43% ($p = 0.08$), adding industrial zones +51% ($p = 0.07$), and the commercial core alone +43% ($p = 0.14$). It is smaller, and not significant, in logs (+36%, $p = 0.23$) and on all parcels including the redeveloped ones (+26%, $p = 0.29$). Residential land values do not rise (−25%, $p = 0.31$), consistent with an office-anchored treatment that operates on commercial land. The weak point is the donor screens. Keeping the centres that received a new Metro station cuts the gap to +16% ($p = 0.22$), and dropping all the screens, which re-admits North Sydney, St Leonards and Liverpool, gives +20% ($p = 0.15$). Those centres are screened because they received large place-based shocks of their own after 2016 (a new station, or their own growth in public jobs), but the size of the headline depends on that exclusion.

5.4.2 Office approvals

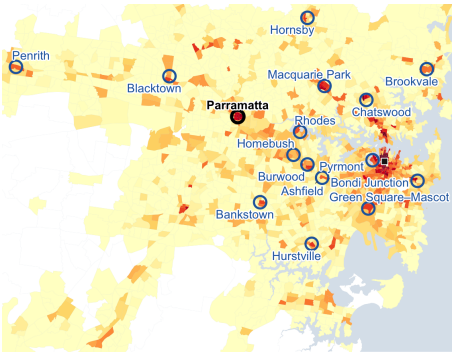
The second outcome is the value of private office construction approved. Parramatta's ABS region (about 8.5 km^2, with 21,994 private office jobs in 2016) approved $1.597 billion of office work over FY2016–17 to FY2019–20, in FY2016–17 prices (deflated by the ABS producer price index for NSW non-residential building construction). The ABS codes a building's owner, not its tenant, so the government-leased towers count as private approvals. I therefore net them out (Figure 14):

- $93m that the ABS itself codes as public;
- $565m for the government floors of the three leased towers: 105 Phillip Street in full, 4 Parramatta Square at a government floor share of 0.96, and 6 and 8 Parramatta Square at 0.50, whose other half is private and stays in the outcome;
- $140m for the original design of 6 Parramatta Square, approved in July 2017 and replaced in 2019.

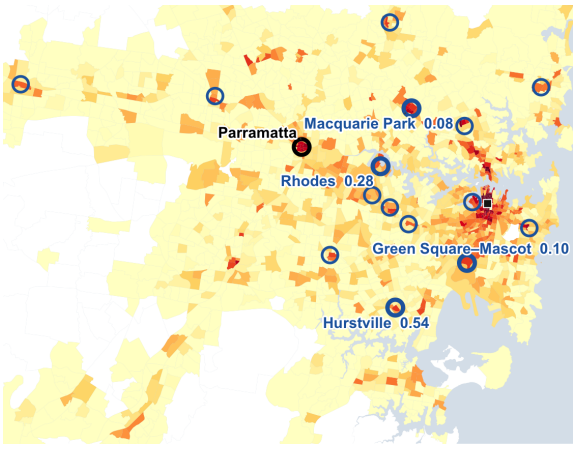
That leaves $799m of private approvals (over 1 billion in 2026 prices): $260m of private floor space inside the government-anchored towers and $539m of standalone private building. The outcome is cumulative real private approvals since FY2012 per 2016 private office job, which puts regions of different size on one scale.



(a) Candidate centres



(b) The 15 most similar, after the screens



(c) Synthetic-control weights

Figure 10: Land-value donor pool: 800 m circles over 2016 jobs per km^2; ■ marks the CBD. (a) Every centre with at least 5,000 jobs within 800 m in 2016, outside the CBD and a 3 km buffer around Parramatta, plus the CBD-fringe office precinct. (b) After dropping single-function centres, centres with a public-jobs shock of their own and centres with a new Metro station opened 2017–25, the 15 most similar to Parramatta on pre-2017 characteristics. (c) The weights from the 2012–16 fit: Hurstville 0.54, Rhodes 0.28, Green Square–Mascot 0.10, Macquarie Park 0.08.

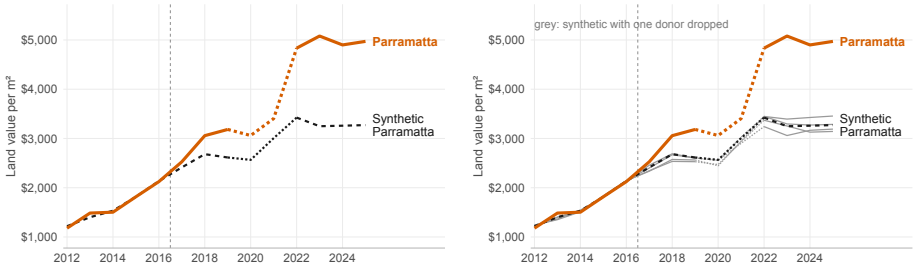


Figure 11: Land value per m^2 of commercially zoned land (excluding industrial), same parcels throughout, 800 m circle: Parramatta against its synthetic control. Dashed vertical: the December 2016 lease. Right: grey lines rebuild the synthetic with each weighted donor dropped in turn (2025 gap +$1,516 to +$1,829). 2019–22 is drawn dotted because the 2020 and 2021 valuations were set by rule (the Valuer General's COVID-19 reduction), not from sales.

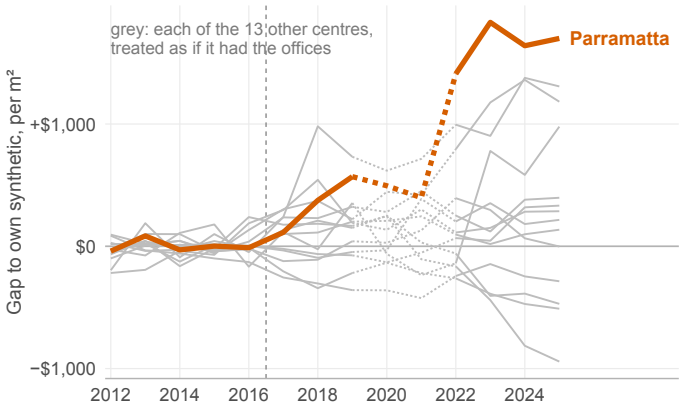


Figure 12: Gap between each unit and its own synthetic control, $ per m^2. Grey: the 13 donor centres whose placebo fits its own pre-period within five times Parramatta's error, each treated as if it had received the offices. Parramatta's post-2017 gap is the largest: $p = 1/14 = 0.07$.

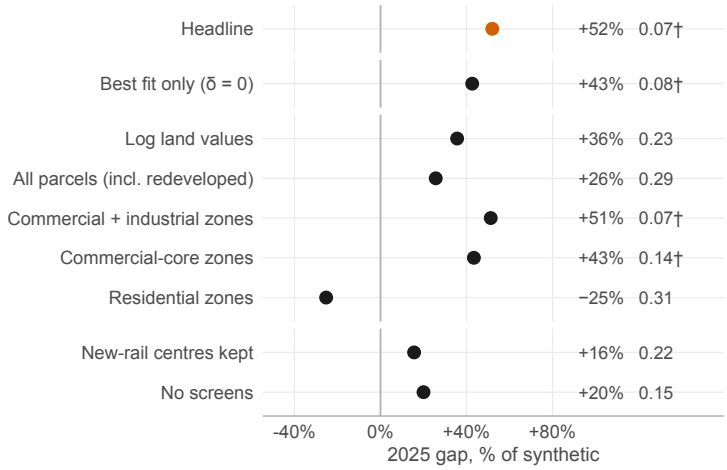


Figure 13: The land-value gap in 2025 under variants of the headline, in per cent of the synthetic level. Right: gap and placebo p (5× pre-fit filter); † marks rows where Parramatta's gap is the largest in the pool. Rows: the headline; the best fit without the similarity tie-break ($\delta = 0$); other outcomes (logs, all parcels including redeveloped ones, other zoning definitions); and the donor pool without the new-rail screen and without any screens.

Figure 15 shows the donor pool. Of the 41 regions containing an employment centre, two border Parramatta and are dropped. Of the remaining 39, 17 approved at least $20m of private offices over FY2012–16, so they had an office pipeline before the policy. Two shock screens then remove regions with a public-jobs change of their own and regions with publicly built offices of their own, leaving 13. The donor pool is the 10 of these most similar to Parramatta on eight pre-2017 characteristics: private job density, the office and knowledge-sector shares, high-rise density, the industrial share, commercial land value, rail access and distance to the CBD. The weights follow the same rule as for land values, fitted on FY2012–16, and three regions carry them: Burwood (0.44), Chatswood (0.29) and Bankstown (0.27). Because the outcome is cumulative, the estimate is the gap that opens over the approval wave, measured from FY2016 to FY2020:

$$\hat{\alpha} = \left(Y_{P,2020} - Y_{2020}^{\text{synth}}\right) - \left(Y_{P,2016} - Y_{2016}^{\text{synth}}\right),$$

multiplied by Parramatta's 21,994 private office jobs to return to dollars.

Over FY2017–20 Parramatta's private approvals exceed its synthetic control's by +$703m (Figure 16). In FY2025–26 construction prices this is about $1.05bn. The wave arrives in FY2019 (3 Parramatta Square and 32 Smith Street) and FY2020 (6 and 8 Parramatta Square). Dropping each weighted donor in turn keeps the gap between +$668m and +$723m. All 10 placebos fit their own pre-period within five times Parramatta's error, and none reaches Parramatta's gap (Figure 17), so $p = 1/11 = 0.09$, the smallest value attainable

with 10 donors. The pre-treatment fit is poor: its error is 38% of Parramatta's pre-period level. Parramatta's pre-period path is essentially one step, the university building at 1 Parramatta Square in FY2015, which no blend of donors traces. The estimate is therefore best read descriptively, as the size of the wave against a small counterfactual. Parramatta's own net private approvals over FY2017–20 are $799m, while every counterfactual puts them at only $88m to $176m. Per private office job, Parramatta's FY2017–20 approvals rank second of 193 NSW regions. After the wave the gap stops growing: FY2021–25 adds only +$38m in this pool.

Figure 18 collects the alternatives:

- Every donor pool, with or without the similarity cut, gives +$660m to +$703m.
- On the strict pool, every weighting rule, onset date, series and netting variant gives +$493m to +$745m. The strict pool also screens out regions with a committed office-supply shock of their own, such as a metro overstation tower, and keeps all ten survivors.
- The two naive benchmarks give +$633m and +$711m: Parramatta's own FY2012–16 approval rate continued, and an equal-weight difference against the pool.

The largest source of uncertainty is the netting of individual approval months to the government towers. The gap is otherwise driven mostly by Parramatta's own wave rather than by the choice of synthetic. After FY2020 nothing is robust: the FY2021–25 gap is small, varies across rows and is negative on the new-build proxy (months of at least $20m). The approval wave ended with 6 and 8 Parramatta Square.

The dollar gap can be converted into office capacity, measured in jobs:

$$\text{jobs} \;=\; \frac{\text{gap}}{\text{overshoot} \times \text{cost per m}^2 \text{ of GFA}} \times \frac{\text{NLA}}{\text{GFA}} \times \frac{1}{\text{m}^2 \text{ per worker}}.$$

Construction cost is $3,732 per m^2 of gross floor area (GFA; range $3,101–$4,296), the GFA-weighted average of five Parramatta towers. Approval values exceed the value of the consented floor space, because they also cover fit-outs and repeat approvals. This overshoot is 1.194 (1.148–1.40), measured on Parramatta's own FY2017–20 private approvals against their consented GFA. Net lettable area, the floor a tenant rents, is 0.77 of GFA (0.68–0.86), the usual band for office buildings. Each worker occupies 17.7 m^2 of lettable floor (14.5–20.7), from the City of Sydney Floor Space and Employment Survey. Together these give 9.8 jobs per $1m of approvals (5.5–13.7). The headline gap of +$703m is therefore room for about 6,900 office workers (3,800–9,600), and the strict pool's +$674m for about 6,600 (3,700–9,200). This is capacity at CBD-normal occupancy, not realised employment: approvals are consents, not built floor space. Most of it was approved in FY2019–20 and delivered after census night 2021, so the employment results below cannot yet see it, and the 2026 census will be the test.

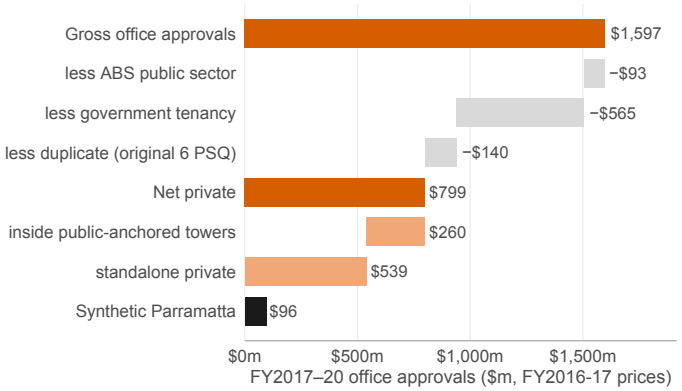
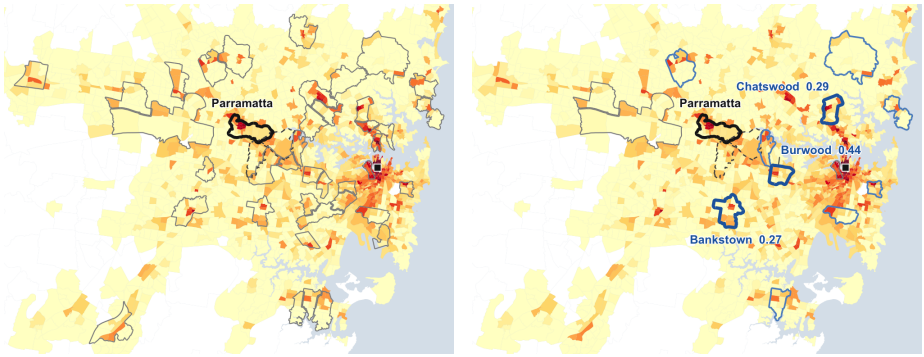


Figure 14: Parramatta's office approvals, FY2017–20, $m at FY2016–17 prices: from the ABS gross total to private approvals. Tenancy: the three government-leased towers at their government floor shares. Duplicate: the first design of 6 Parramatta Square, replaced in 2019. The final bar is the synthetic control's approvals over the same years ($96m).



(a) The 39 candidate regions (b) The 10 most similar and their weights

Figure 15: Office-approvals donor pool: ABS regions over 2016 jobs per km^2; dashed, the 3 km buffer; ■ marks the CBD. (b) Blue: the 10 regions most similar to Parramatta among those passing the data and shock screens (bold: carrying weight, Burwood 0.44, Chatswood 0.29, Bankstown 0.27); grey: the other three passing the screens.

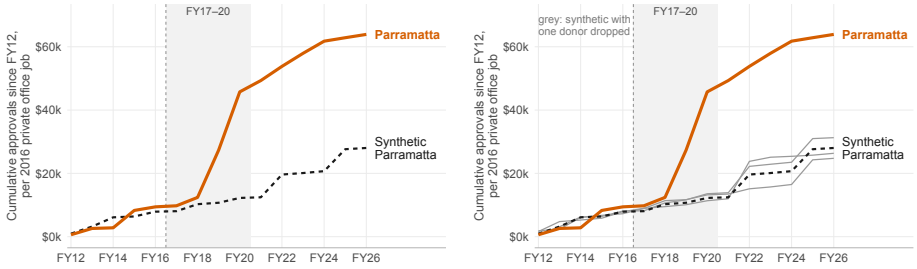


Figure 16: Cumulative real private office approvals since FY2012 per 2016 private office job: Parramatta against its synthetic control (narrow pool, 10 most similar). Shaded: the FY2017–20 wave. Pre-period error 38% of Parramatta's level. Right: grey lines drop each weighted donor in turn (FY2017–20 gap +$668m to +$723m).

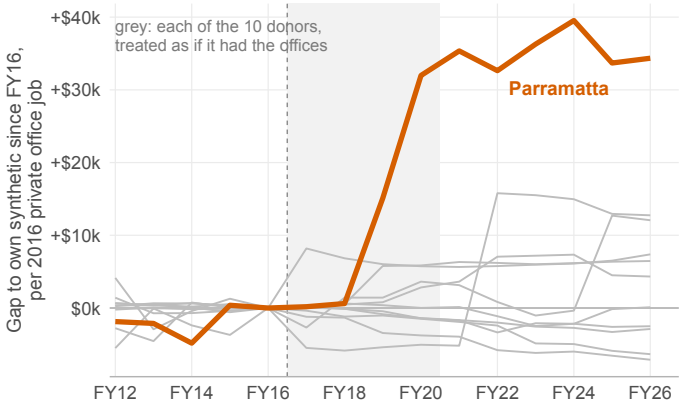


Figure 17: Gap to own synthetic control, re-based at FY2016. Grey: each donor treated as if it were Parramatta, on its own donor pool fixed in advance. $p = 1/11 = 0.09$, the floor.

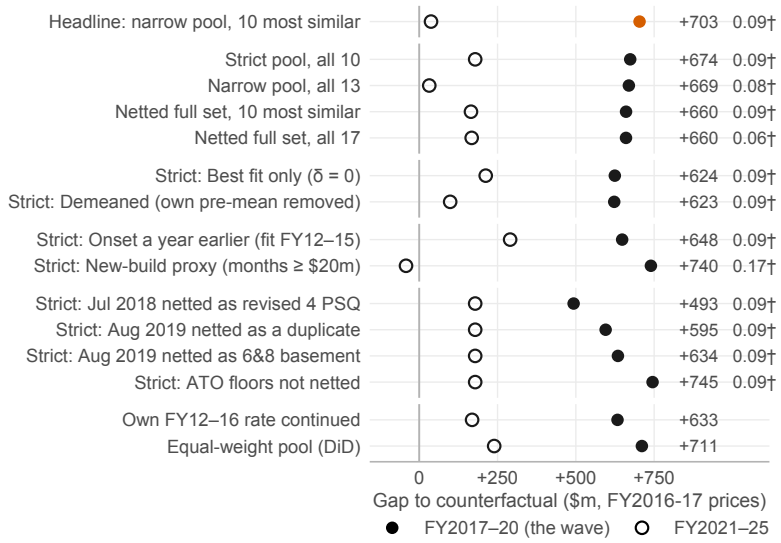


Figure 18: The office-approvals gap under alternative pools, rules, series, netting readings and naive benchmarks, $m at FY2016–17 prices. Filled: FY2017–20 (the wave); hollow: FY2021–25. Right: the FY2017–20 gap and placebo p; † marks rows where Parramatta's gap is the largest. The rule, timing, series and netting rows vary the strict pool, the only design they were run on.

5.4.3 employment

As stated above, it is too early to assess the overall impact on employment in Parramatta owing to the fact that the last reliable employment estimates come from the 2021 census, which occured during lockdowns and before a large number of civil servants had moved into their new Parramatta offices. However, comparing employment growth rates in Parramatta and the comparison units, it appears that up 2021:

- Public employment increased substantially: even before the new post 2021 jobs arrived, government employment in the 800 m circle rose from 13,003 in 2011 to 22,812 in 2021, a gain of 9,809 jobs equal to 27% of the circle's total 2011 employment of 36,230, and the government share of employment in the circle rose from 36% to 48%. NSW state government – the level of government the policy actually moves – accounts for almost all of it (9,120 to 18,759, +9,639). The increase is a step and not a trend: government employment is flat between the 2011 and 2016 censuses (+126) and then jumps by 9,684 over 2016–2021, of which state government is +8,548. Figure 35 plots the series against the documented office moves, with the level-of-government and industry decompositions in Figure 20.

- Overall private employment appears relatively flat, see Figure 36: private employment in the circle goes 23,498 (2011) to 27,686 (2016) to 25,497 (2021), a 2011–2021 multiple of ×1.085 that ranks 13th of the 24 units in pool V1 – squarely mid-pack among the comparison units, and on the same lockdown-censored 2021 endpoint as the public series. There is no sign of a broad private-sector expansion accompanying the office moves.

- There may have been an increase in ANZSIC division M: professional, scientific and technical services. Division-M employment in the circle grows from 2,752 (2011) to 4,388 (2021), ×1.59, the fastest growth of the 24 units in pool V1 (Figure 37). Government employers account for only a small part of it (government jobs within M grow 72 to 282): netting them out the multiple is ×1.53 and the rank 3rd, behind only two centres growing off bases under 400 jobs, while the Sydney CBD sits at ×1.47. Parramatta is therefore unremarkable in private employment overall and first in professional and technical services.

The same synthetic-control design, applied to census employment, finds no detectable private response by 2021 (Figure 19). The outcome is private business employment by place of work in the 800 m circles, observed at the five censuses 1996–2016 before treatment and in 2021 after it. The pool comes from the same screens and the same top-15 rule as for land values, except that the new-rail screen covers only stations opened before the 2021 census, so North Sydney (whose Metro station opened in 2024) stays in. The weights use the same $\delta = 1$ tie-break, with similarity measured on 2011 industry composition and resident density. Relative to its synthetic control, Parramatta's 2021 private business employment is 29% lower ($p = 0.17$). Retail and hospitality is 5% lower ($p =$

0.88), professional, scientific and technical services 4% higher ($p = 1.00$) and knowledge services 29% lower ($p = 0.10$). Matching on the best pre-treatment fit alone gives the same picture. The negative point estimates come from the donors rather than from Parramatta: North Sydney's employment boomed over 2016–21, and it carries 0.56 of the weight in the knowledge-services synthetic. ABS counts of businesses up to 2025, which exclude general government, likewise show no significant change in any industry group, whether for all businesses or for those with five or more employees. None of this is yet a test of the policy's private response. The 2021 census was taken in lockdown and before 6 Parramatta Square opened, and most of the office capacity approved in the wave was delivered after census night. The 2026 census will provide a more clean test.

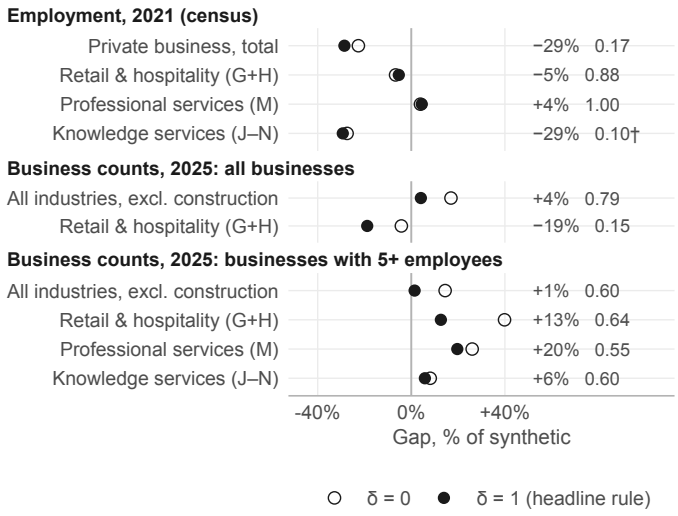


Figure 19: Synthetic-control gaps for private employment and business counts, in per cent of the synthetic level. Top: census employment at place of work, 800 m circles, 2021 (a lockdown census, taken before 6 Parramatta Square opened). Middle and bottom: ABS counts of businesses (which exclude general government) in 2025, all businesses and those with five or more employees. Filled: the headline $\delta = 1$ rule; hollow: the best pre-treatment fit alone ($\delta = 0$). Right: the $\delta = 1$ gap and placebo p; † marks rows where Parramatta's gap is the most extreme in the pool.

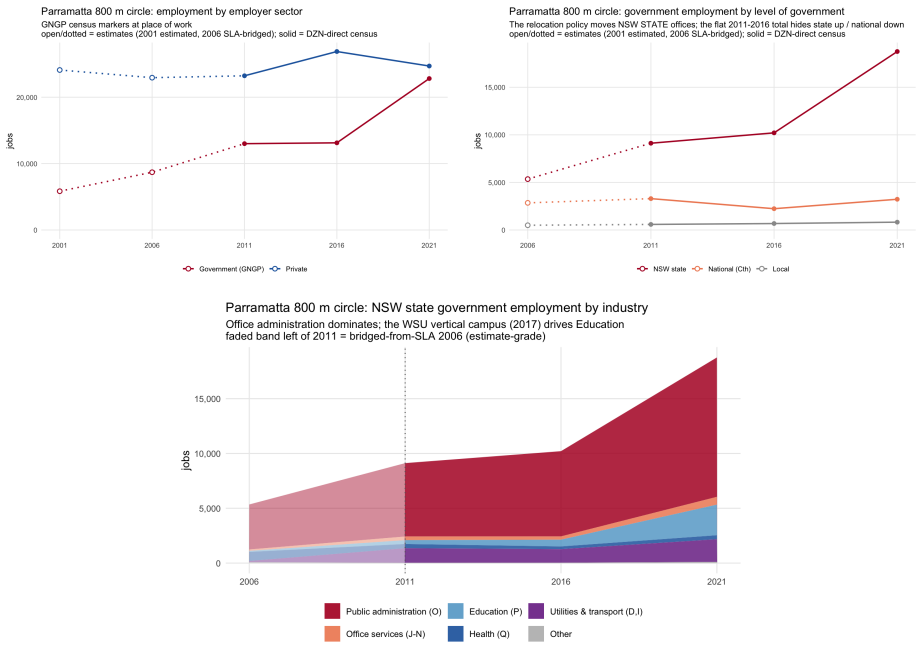


Figure 20: Public and private employment in the Parramatta 800 m circle (census GNGP employer markers, place of work). Top left: government vs private employment. Top right: government employment by level of government – the 2011–2016 plateau in total government conceals rising NSW state employment offset by falling Commonwealth employment. Bottom: NSW state government employment by industry; office administration dominates throughout, with the WSU vertical campus (2017, 1 PSQ) driving the Education slice. Open markers and dotted segments mark the estimate-grade years – 2001 is estimated from the industry mix and 2006 is bridged from SLA-level data – while solid segments are direct Destination-Zone measurements; the faded band in the bottom panel marks the same estimate-grade era.

6 Public sector shift in the model

The reduced form results give us a sense of the local impacts of the policy, but cannot capture general equilibrium effects on the rest of the city or overall welfare. I therefore extend my quantitative urban model to include a public sector so that I can study the impacts of the policy in the model equilibrium. The public sector is separate from the private sectors, and is controlled by a policy maker that dictates where to rent offices for its workers. The workers' residences are fixed in place, and their welfare is not considered in counterfactuals. The key model parameter is the spillover from public to private sector firms: if these spillovers are high, the government can pull private sector jobs across the city just by moving its own workers. Moving public sector jobs also impacts the private sector by changing local floorspace use and rental prices.

Having calibrated the model to 2016 data, I find the counterfactual equilibrium in which public sector workers are moved to Parramatta, and by comparing the results to the reduced form synthetic controls results, I test for the implied public-private spillovers, and thereby the welfare impacts of the policy. The census data shows that the number of public sector workers in Parramatta increased by 9,684 from 2016-2021. A further 4200 public sector workers were scheduled to occopy 6 Parramatta square after its opening in 2022. I therefore move 13,884 public sector jobs in the counterfactual.

The effect depends crucially on the value of the public to private spillovers. In Figure 21, I show how the steady state output, high type welfare and low type welfare depend on the public to private spillover weight; results are presented both for the case where high and low types have the same private to private spillovers, and the case in which only high type private workers provide productivty spillovers due to agglomeration.

When spillovers are high, moving public sector jobs leads to higher employment in their new location and lower employment around the CBD. Because the CBD is the highest productivity location, this decentralisation leads overall output to fall. Because the CBD is closer to where graduates live and uses a more graduate intensive production technology, graduates lose from this dispersion; non graduates live closer to the new hub and the production technology there is less graduate intensive, so non graduates gain from this dispersion. When public-private spillovers are low, however, moving public sector jobs out of the CBD just frees up space for private sector workers there and reduces spaces where they relocate to, so the effects switch: workers centralise more in the high productivity, graduate intensive CBD, and output and graduate welfare go up while non-graduates lose.

6.1 calibrating the spillovers to the data

I provide a rough calibration of the spillovers to the observed shocks in the data. Given that we only have access to sufficent post period data on land values and office approvals data, I test the model calibration to the change in land values and the change in private employment implied by the private office approvals.

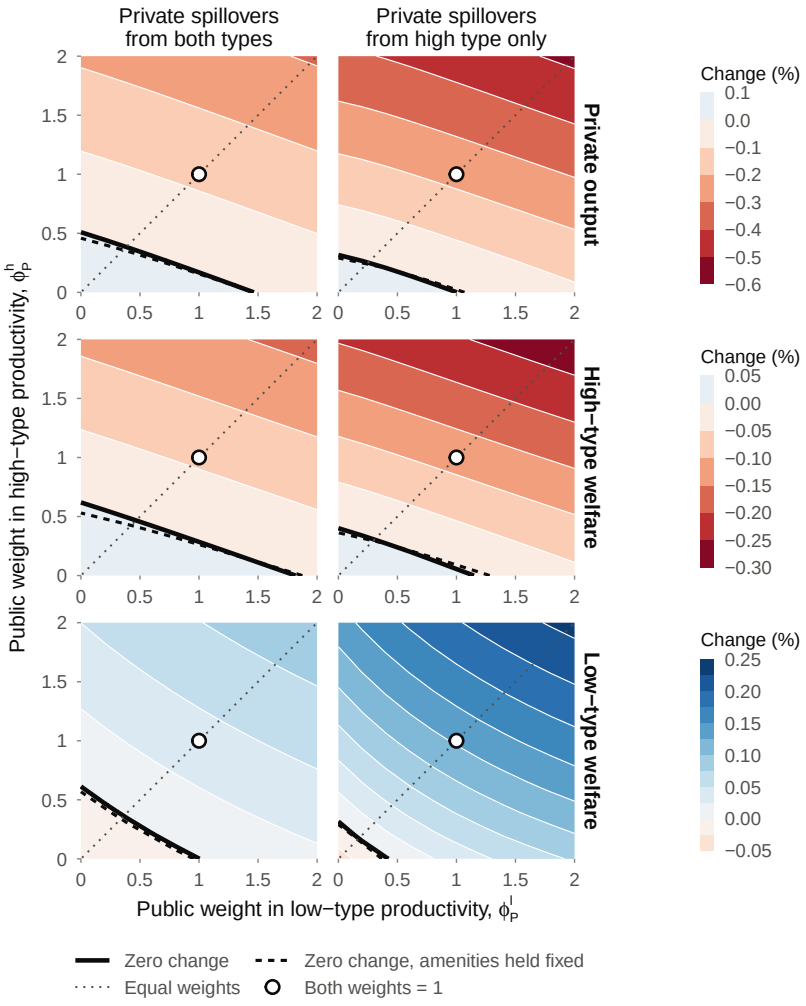


Figure 21: Effects of the relocation by public spillover weight, static model with a CES labour bundle ($\sigma = 1.4$). Each panel shows the percentage change from the 2016 baseline equilibrium when 13,884 public jobs are moved to Parramatta at once, with the quantity of public office floorspace fixed. ϕ_P^h and ϕ_P^l are the weights on public employment in the productivity spillovers to high- and low-type labour (ϕ_P in equation (59), here allowed to differ by receiving type). Left column: private spillovers come from private workers of both types ($\phi_l = 1$); right column: from high-type private workers only ($\phi_l = 0$). The model is recalibrated to the same 2016 data at every point. Output is private output and welfare is that of private workers. The dashed line is the zero line with residential amenities held at their baseline values; the dotted line marks equal weights, and the circle weights a public job exactly like a private one. Solved on a 0.2 grid (0.4 with amenities fixed), interpolated bilinearly.

In the static model, land value is proportional to floorspace expenditure. Expenditure on floorspace at i is residents' housing spending plus office spending, which under Cobb–Douglas production is $c_\alpha \equiv (1-\alpha)/\alpha$ times the wage bill (7), plus the floorspace rented by public offices:

$$X_i \equiv Q_i H_i = \sum_t (1-\beta^t) \sum_j w_j^t L_{ij}^t + c_\alpha \sum_t w_i^t L_{Fi}^t + Q_i H_{Pi}. \tag{34}$$

Floorspace supply $H_i = K_i \psi_i Q_i^\zeta$ gives $X_i = K_i \psi_i Q_i^{1+\zeta}$, the same power of the rent as the land value in (21), so

$$K_i V_i = \frac{X_i}{r(1+\zeta)}, \qquad \text{hence} \qquad \mathrm{d}\ln V_i = \mathrm{d}\ln X_i, \qquad \mathrm{d}\ln Q_i = \frac{\mathrm{d}\ln X_i}{1+\zeta}, \qquad \mathrm{d}\ln H_i = \frac{\zeta\,\mathrm{d}\ln X_i}{1+\zeta}. \tag{35}$$

A change in land values is therefore a change in floorspace expenditure, one for one, whatever the cost shifter ψ_i: a high building cost means a high rent and little floorspace, and the two cancel in the product. ζ only splits the change between rent and floorspace.

To see what moves floorspace expenditure, take a small location j with rents and wages elsewhere held fixed, and let s_j be the share of X_j spent by private offices, the rest by residents. Zero profit (48) sets the wage, jobs follow the wage with elasticity ε, and residents leave as the rent rises with elasticity $\varepsilon(1-\beta)$:

$$\mathrm{d}\ln w_j = \tfrac{1}{\alpha}\,\mathrm{d}\ln A_j - c_\alpha\,\mathrm{d}\ln Q_j, \qquad \mathrm{d}\ln L_{Fj} = \varepsilon\,\mathrm{d}\ln w_j, \qquad \mathrm{d}\ln L_{Rj} = -\varepsilon(1-\beta)\,\mathrm{d}\ln Q_j.$$

With $\mathrm{d}\ln Q_j = \mathrm{d}\ln X_j/(1+\zeta)$ from (35), land value responds to local productivity, and to new public office spending x_j^P (as a share of X_j), as

$$\mathrm{d}\ln V_j = \frac{1}{D_j}\Big[\frac{s_j(1+\varepsilon)}{\alpha}\,\mathrm{d}\ln A_j + x_j^P\Big], \qquad D_j \equiv 1 + \frac{s_j(1+\varepsilon)\,c_\alpha + (1-s_j)\,\varepsilon(1-\beta)}{1+\zeta}. \tag{36}$$

The numerator is the rise in office spending at the old rent, the wage plus the jobs that follow it; $D_j > 1$ is the damping as the rent rise cuts wages, and so jobs, and prices residents out. Because A_j is Hicks-neutral, both types' wages move in proportion and σ drops out to first order; $1-\beta$ is the spending-weighted mean of the two types' housing shares $1-\beta^t$. At the static parameters ($\alpha = 0.8$, $\zeta = 1.86$, $\varepsilon = 6.2$), a location that is all offices ($s_j = 1$) has

$$\frac{\mathrm{d}\ln V_j}{\mathrm{d}\ln A_j} = \frac{(1+\varepsilon)/\alpha}{D_j} = \frac{9.0}{1.63} = 5.5,$$

and Table 5 works through a one per cent rise in productivity. Land value amplifies productivity: jobs follow the wage, and builders house them.

6.1.1 Results

The value of the local productivity change that explains the reduced form change in LV shows that local productivity must increase by 5.6%, while it must increase

Productivity A_j +1 %		
Wage	+0.77 %	$1/\alpha = 1.25$ % at the old rent, less $c_\alpha\times$ the rent rise
Jobs	+4.76 %	$\varepsilon\times$ wage: workers follow the wage
Office spending = land value	**+5.52 %**	wage + jobs
Rent	+1.93 %	$\div(1+\zeta)$
Floorspace	+3.59 %	$\times\zeta/(1+\zeta)$: builders absorb most of it

Table 5: Response of a location that is all offices ($s_j = 1$) to a one per cent rise in productivity, from (35) and (36) at $\alpha = 0.8$, $\zeta = 1.86$, $\varepsilon = 6.2$, with rents and wages elsewhere held fixed. The numerator $(1+\varepsilon)/\alpha = 9.0$ is the response with no rent feedback; $D_j = 1.63$ is the damping from the rent rise.

by 2.6% to fill only the new private office space implied by the approvals series up to 2026. In both cases, the public-private spillover parameter required to explain the model is very large (¿5), as it reaches the limits of what the current specification of public-private spillovers can explain. Taking it literally, they imply that we are in the realm where overall output falls along with graduate welfare, while non-graduate welfare rises. In future I intend to test alternate specifications in which the elasticity λ is allowed to vary by type, or in which I try to capture effects on local construction productivity and to net out the effects of the tramline.

References

[1] Alberto Abadie, Alexis Diamond, and Jens Hainmueller. Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program. *Journal of the American statistical Association*, 105(490):493–505, 2010.

[2] Alberto Abadie and Jérémy L'Hour. A penalized synthetic control estimator for disaggregated data. *Journal of the American Statistical Association*, 116(536):1817–1834, 2021.

[3] Gabriel M Ahlfeldt, Stephen J Redding, Daniel M Sturm, and Nikolaus Wolf. The economics of density: Evidence from the berlin wall. *Econometrica*, 83(6):2127–2189, 2015.

[4] Nathaniel Baum-Snow and Lu Han. The microgeography of housing supply. *Journal of Political Economy*, 132(6):1897–1946, 2024.

[5] Sascha O. Becker, Stephan Heblich, and Daniel M. Sturm. The impact of public employment: Evidence from Bonn. *Journal of Urban Economics*, 122:103291, 2021.

[6] Pierre-Philippe Combes, Gilles Duranton, and Laurent Gobillon. The production function for housing: Evidence from France. *Journal of Political Economy*, 129(10):2766–2816, 2021.

[7] Giulia Faggio. Relocation of public sector workers: Evaluating a place-based policy. *Journal of Urban Economics*, 111:53–75, 2019.

[8] Giulia Faggio, Teresa Schlüter, and Philipp vom Berge. Interaction of public and private employment: Evidence from a German government move. *Regional Science and Urban Economics*, 111:104084, 2025.

[9] Dimitria Freitas. The potential of public employment reallocation as a place-based policy. NBER Working Paper 33432, National Bureau of Economic Research, 2025.

[10] Edward L Glaeser and Matthew E Kahn. Sprawl and urban growth. In *Handbook of regional and urban economics*, volume 4, pages 2481–2527. Elsevier, 2004.

[11] Mariaflavia Harari. Cities in bad shape: Urban geometry in india. *American Economic Review*, 110(8):2377–2421, 2020.

[12] Stephan Heblich, Stephen J Redding, and Daniel M Sturm. The making of the modern metropolis: evidence from london. *The Quarterly Journal of Economics*, 135(4):2059–2133, 2020.

[13] Mi Seon Jeon and Jong Soo Lee. Is public agency relocation effective to achieve decentralization? evaluating its effects on regional employment. *Journal of Urban Affairs*, 45(8):1486–1501, 2023.

[14] Hans RA Koster. The welfare effects of greenbelt policy: Evidence from england. *The Economic Journal*, 134(657):363–401, 2024.

[15] Max Nathan, Henry Overman, Capucine Riom, and Maria Sanchez-Vidal. Multipliers from a major public sector relocation: The BBC moves to Salford. IZA Discussion Paper 17337, IZA Institute of Labor Economics, 2024.

[16] John Nicholas Tsivanidis. The aggregate and distributional effects of urban transit infrastructure: Evidence from bogotá's transmilenio. 2018.

A Additional static counterfactuals

This appendix reports the stylised counterfactuals that complement the CBD shift of Section 4.1: sewing the harbour together, compacting the city inward to fill the harbour, and compacting it to a circle.

A.1 Sewing the harbor together

The CBD's eastern location means that its average distance from other points in the city is larger than if it were more centrally located. However, its accessibility to neighbouring areas is more immediately impacted by the difficulty of traversing the harbor itself. Indeed, before further investigation it is possible that alleviating this local accesibility barrier is what drives the gains in the above cbd shift counterfactual. To understand the costs of the harbor travel barrier, I consider a counterfactual equilibrium in which the harbor is 'sewn together': adjacent sides of the harbor are identified and travel times between them set to zero. Very roughly speaking, if we consider Sydney as a circle with a pizza slice cut out (the harbor), the counterfactual identifies the two sides of the empty slice (geometrically equivalent to folding up the circle into a cone). [19]

I implement the sewing in two stages:

1. Cross harbor links: My aim is to acheive something close to the metaphor of overcoming the harbor by folding up sydney into a cone, such that only directly opposite sides of the harbor are identified. If we were to allow travel times from the far south east to far north east of the harbor to fall to zero, this would represent a more drastic change than this slice traversing analogy, and provide far more favorable local geometry than for a city set on a circle of equal size. Therefore, to identify only roughly opposite sides of the harbor, I first draw a line through the center of the harbor. I acheive this by computing the ridge of points equidistant from the two shores and extract from it three smoothed centreline arms — the main channel from the upper Parramatta River to the Heads (27.3 km), Middle Harbour (13.1 km) and the Lane Cove River (9.0 km). Along these centrelines, I create waypoints every 400m. For each waypoint, I then find the closest occupied hexagons on either side of the centreline, and link them at their road-network anchor points. To prevent non-physical pathologies, I cap the number of links such that for each hexagon the sum of links and existing adjacent occupied (ie not harbor) hexagons is at most 6 and no links cross eachother. Of the 117 waypoints, 67 forge a link. See Figure 22 for the resulting linkages.

[19] One might worry that transforming to cone geometry does more than just restoring travel times across the empty slice. The effects are small: computing mean pairwise geodesic distances by simulation, a cone folded from a disc with a slice of up to 30° removed differs from the flat disc of equal area by less than 0.1% (and by less than 0.5% even at 60°), so this aspect of the counterfactual shouldn;t be driving results.

2. Imputing travel times: First, I estimate the average travel time between neighboring inner city hexagons. Precisely, I take the set of neighboring inner city hexagons and estimate the following function f(d)= alpha+ beta*d for the travel time between the R5R snapped station points of those pairs according to the R5R travel time algorithm ($f(d) = 1.72 + 1.85\, d_{\text{km}}$ minutes at free flow). I then set the travel time between linked harbor hexes equal to the time taken to travel from the first hex anchor to its shoreline (along the chord between the two hex anchors), instantaneously jump across the harbor, then travel from the linked harbor hexes shoreline to its center (ie I sum the inferred travel time from each of their anchors to their own nearest shoreline according to this f(d)). This gives me the travel time between linked hexes, but then I also need to update the rest of the travel time matrix so that journeys between other hexes can make use of the new links– eg from hexes one row above the north harbor shoreline to those one row below the south harbor shoreline. I implement this using a min–plus shortest-path update. First, I assume any journey through the new links will take the new path from the centroid- shorline and shoreline-centroid of newly linked hexes (ie these are acting as bridge entrances). Then I update the matrix of travel times between the whole set of hexes which have new habor front links with a Floyd–Warshall algorithm. Then, as the journey through between other origin-destination pair can only be affected by taking a journey through these newly linked hexes, I am able to update every other origin–destination pair as $\tau'_{ij} = \min\big(\tau_{ij},\, \min_{(a,b)}(\tau_{ia} + c_{ab} + \tau_{bj})\big)$. The result is that cross hex pairs have reduced travel times, with a maximum reduction of 41.1 minutes; 2.2% of all origin–destination pairs improve, by 2.9 minutes on average, and no pair worsens. See Figure 23 for the map of travel time reductions and Figure 24 for the distribution of travel time reductions across the improved pairs.

A.1.1 Results

In the baseline specification of the basic model, sewing the harbor increases output by 0.6%, commutes fall by 0.6% (from 20.96 to 20.83 minutes), and welfare increases by 0.34%.[20] Total rents increase mechanically by the same amount as output 0.6%, with rents rising near the harbor and falling elsewhere. The number of people in the highest employment 400m hexagon increases by 1.8% (from 103,000 to 105,000), and as shown in Figure 25, the employment share of the highest n employment hexagons is higher in the new counterfactual for all n — there is no crossing point. This is partly driven by more residents concentrating

[20] As throughout, the headline uses the empirically anchored crossing price and live amenities, and should be read with its brackets: pricing the crossings with the global land-trip fixed term instead gives +0.14%, and retaining the water width and pricing it as ordinary land gives +0.07%; holding residence amenities fixed at baseline the headline becomes +0.23%. Sewing shut only the 36-cell patch of the inner harbor (the same footprint the fill counterfactual of the next section develops) is worth +0.35% — almost the whole seam.

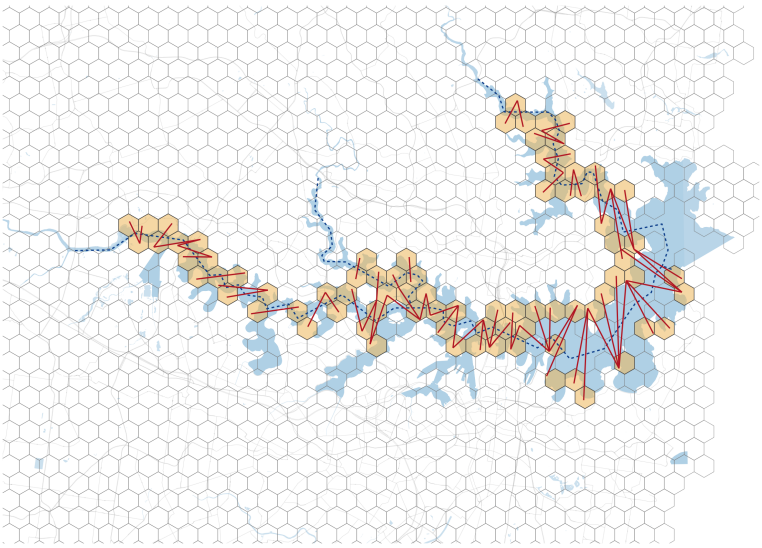


Figure 22: The harbor seam: the 67 accepted crossing links (red), drawn between the road-network anchor points of the linked hexagons (shaded orange). The dashed lines are the three centreline arms — the main channel, Middle Harbour and the Lane Cove River — whose 400 m waypoints nominate the links. Each link deletes its over-water width from the travel-time matrix; its land legs remain priced at land cost.

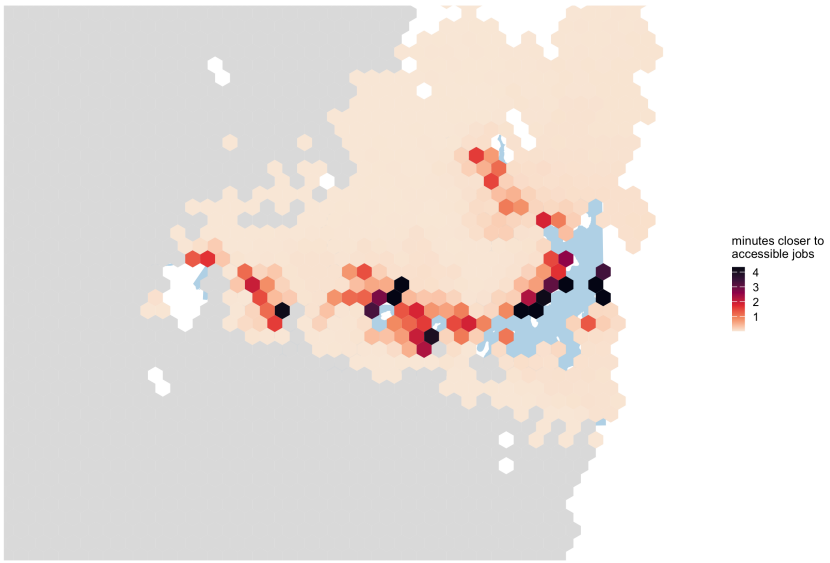


Figure 23: Travel-time reductions from sewing the harbor: for each hexagon, the weighted mean reduction in travel time to jobs, weighting each destination by its employment discounted by commuting cost (weights $L_{Fj}\, e^{-\nu \tau_{ij}}$ at baseline travel times, own hexagon excluded) — "minutes closer to accessible jobs". The window shows the harbor region (33×26 km) of the SUA. Hexagons with reductions below 0.01 minutes are drawn grey; the colour scale is clipped at the 99th percentile (4.3 minutes) — the largest foreshore value is 8.2 minutes. The gains hug the two shores and decay quickly inland.

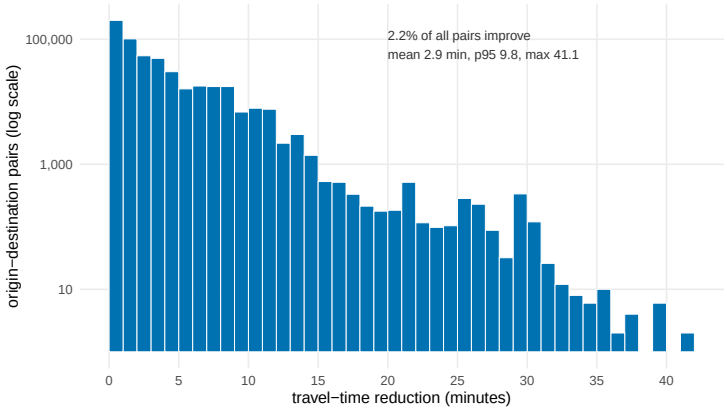


Figure 24: Distribution of travel-time reductions across improved origin–destination pairs (log count scale). The 97.8% of pairs left unchanged by the sewing are excluded from the plot; 2.2% of pairs improve, by 2.9 minutes on average (95th percentile 9.8 minutes, maximum 41.1 minutes).

around the northern side of the harbor, which now has reduced travel times to the areas around the CBD on the south side of the harbor (Figure 27). Of the 41,240 residents who relocate, 22,075 settle in newly densified hexagons on the northern side of the harbor against 19,165 on the southern side; netting each side's thin losses, the north side gains 6,241 residents (and 3,174 jobs) from the south. Cross-harbor commuting rises by 15,467 journeys (+3.2%, from 483,000 to 499,000): north-to-south commutes — residents of the north side working on the south, already the dominant direction at baseline — rise by 9,267 (+2.9%) to 329,000, and south-to-north commutes by 6,200 (+3.8%) to 170,000.[21] Overall the counterfactual results suggest that the inner city transport barrier imposed by the harbor is depressing inner city employment concentration, which in turn reduces wages and output, and adds to the effect of longer commutes in reducing welfare.

The sewing gain of 0.34% is about a quarter of the CBD-shift gain (1.4%), and about a fifth of the next section's harbor fill reshuffle under its conservative brackets (+1.9% with the built floorspace stock moved as fixed, +1.2% with amenities frozen). Thus the model inffered costs of the eastern, cbd location appear to operate mainly through its non-central location, and the the harbors costs appear arise mostly through preventing a compact city shape. Note these

[21] Sides are classified by splitting the city along the main-channel centreline of Figure 22, extended west along the Parramatta River and east past the Heads, so both banks of Middle Harbour and the Lane Cove River count as north; a commute is cross-harbor when residence and workplace lie on opposite sides, and the flows are the model's commuting probabilities summed over those pairs. Restricted to commutes with both ends inside the harbor window of Figure 23 — crossings of the estuary itself — the increases are 12,054 north-to-south (+6.3%) and 7,523 south-to-north (+8.4%): larger than the city-wide changes, because cross-side commuting far upstream falls as the city re-sorts toward the crossing.

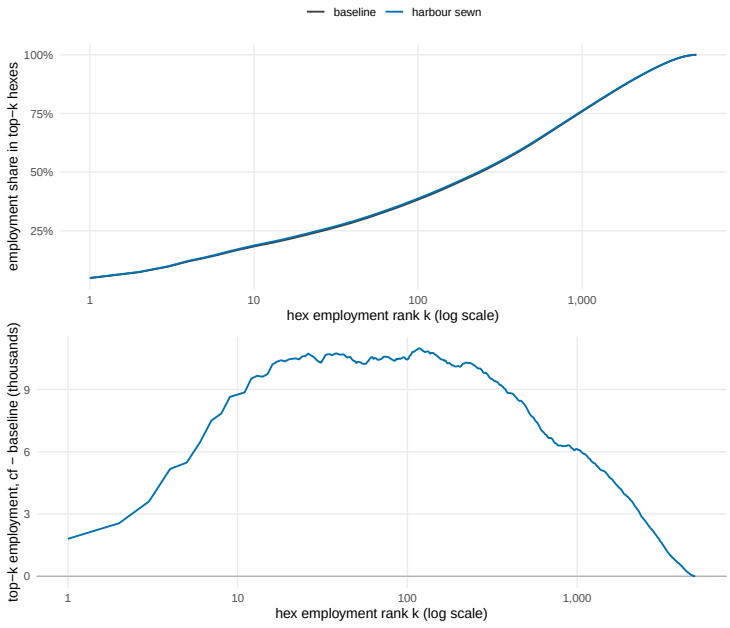


Figure 25: Employment concentration, baseline vs the sewn harbor. Top panel: cumulative employment share of each city's top-k hexagons (each scenario ranked by its own hexagon employment). Bottom panel: the difference between the two curves. The counterfactual lies above the baseline at every rank, with the surplus peaking at 11,000 workers around $k = 117$ — a modest city-wide tightening of concentration, an order of magnitude smaller than under the CBD shift.

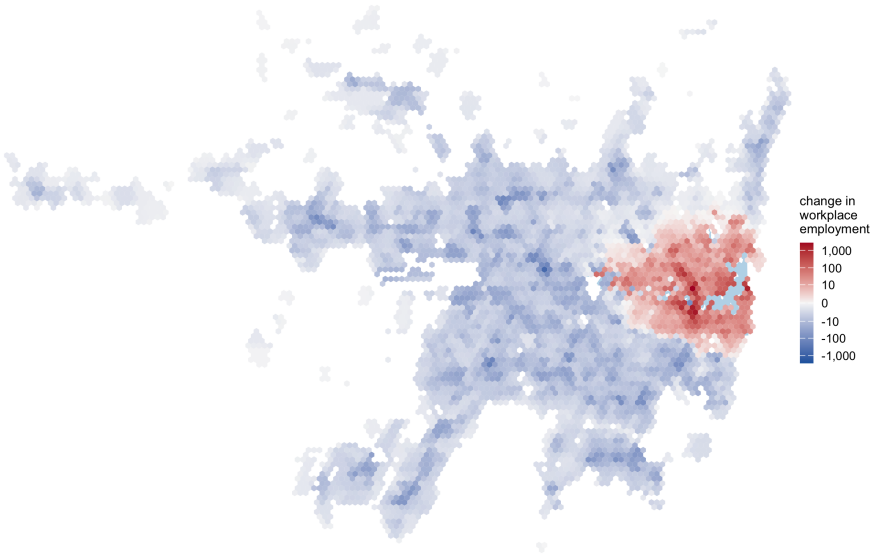


Figure 26: Where employment is gained and lost when the harbor is sewn (headline specification, Sydney SUA, 400 m hexagons, no spatial aggregation). Each hexagon is its change in workplace employment, counterfactual minus baseline, on a signed-logarithmic diverging scale; total employment is fixed, so the map is a redistribution. The gain is tightly localised on both shores of the crossing — every gaining cell lies within the area covered by Figure 23, and the ten largest gainers take 41% of the total gain of 29,633 workers — while the offsetting loss is spread thinly across the rest of the city, nowhere exceeding 637 workers in a single hexagon.

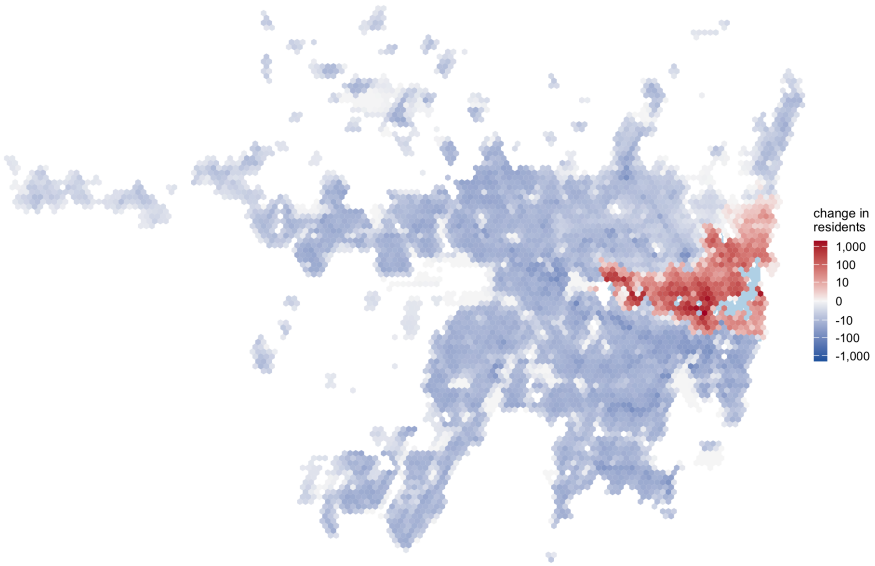


Figure 27: Where residents move when the harbor is sewn (headline specification, Sydney SUA, 400 m hexagons, no spatial aggregation). Each hexagon is its change in residents, counterfactual minus baseline, on a signed-logarithmic diverging scale; total residence is fixed, so the map is a redistribution. In total 41,240 residents relocate — more than the 29,633 workers of Figure 26 — and the gain lands entirely within the area covered by Figure 23, hugging both shores of the sewn crossings. The residential pull is less top-heavy than the employment pull (the ten largest gaining hexagons take 28% of the gain, against 41%), and the offsetting drain is spread so thinly across the rest of the city that no hexagon loses more than 87 residents.

results are preliminary: currently, transport costs do not account for congestion and more checks are needed of the mechanisms driving the counterfactuals.

A.2 Compacting the city

The 'Harbor sew' counterfactual gives an indication of costs of the harbor as a travel barrier, but this is not the only cost it imposes on the city. It also contributes more generally to Sydney's non-compact urban form: the undeveloped harbor contributes to the effect of other physical barriers (the ocean, the mountains) that together force the city to sprawl far to the west and along the coast as it expands. To understand these costs, I run a series of counterfactuals reshuffling the city to fill into currently undeveloped areas.

In the first counterfactual, I simulate 'filling in the harbor'. To make the results more cleanly interpretable, I keep the total number of active hexagons fixed. This helps me to avoid problems that would be introduced by simply allowing development on previously unoccupied harbor cells (such as the problem of how to allocate fundamental productivity or construction for these cells and the fact that worker utility would mechanically increase by allowing more frechet utility draws). Instead, I reshuffle the attributes of existing hexagons to fill the harbor. To minimize the distortion to the city, I make the minimum number of reshuffles necessary to fill the harbor while adhering to the spirit of compacting the city. First, I construct a target footprint of hexagons that I will reshuffle existing hexes onto. I start with the existing set of occupied hexes, then remove the hexes furthest from the CBD one by one until the same number have been removed as are required to fill the harbor. I then shuffle the existing hexagons to fill the new, harbor-inclusive urban footprint while minimizing the sum of squared distances traveled by each individual hexagon. Formally, let $\mathcal{X}$ denote the set of N centroid positions in the original footprint of active hexagons and $\mathcal{Y}$ the set of N positions in the target footprint, and let Σ be the set of bijections $\sigma : \mathcal{X} \to \mathcal{Y}$. I find the bijection of hexagon positions from the original footprint to the new footprint that minimizes the sum of squared per hexagon journey's from origin to finishing point:

$$\min_{\sigma \in \Sigma} \sum_{i=1}^{N} \left\| p(i) - \sigma\big(p(i)\big) \right\|^2 \tag{37}$$

I move hexagon i's entire bundle of attributes to the site $\sigma^\star(p(i))$. This is a discrete optimal transport problem. It is implemented through the Jonker–Volgenant algorithm using MATLAB's built-in `matchpairs` solver. Because the cost is over the journey traveled per hexagon squared, large displacements of a single hexagon from its origin are penalized more than shifting many individual hexagons to make up the same total distance. Thus instead of the harbor hexagons being filled by directly transporting hexes all the way from the emptied fringe, it is filled by shuffling in hexagons at the harbor shore a few positions, then the row behind them shuffling in to take their place and so on. Figure 28

shows the map of the paths taken by the hexagons, and Figure 29 shows the map of distance traveled by each hexagon.

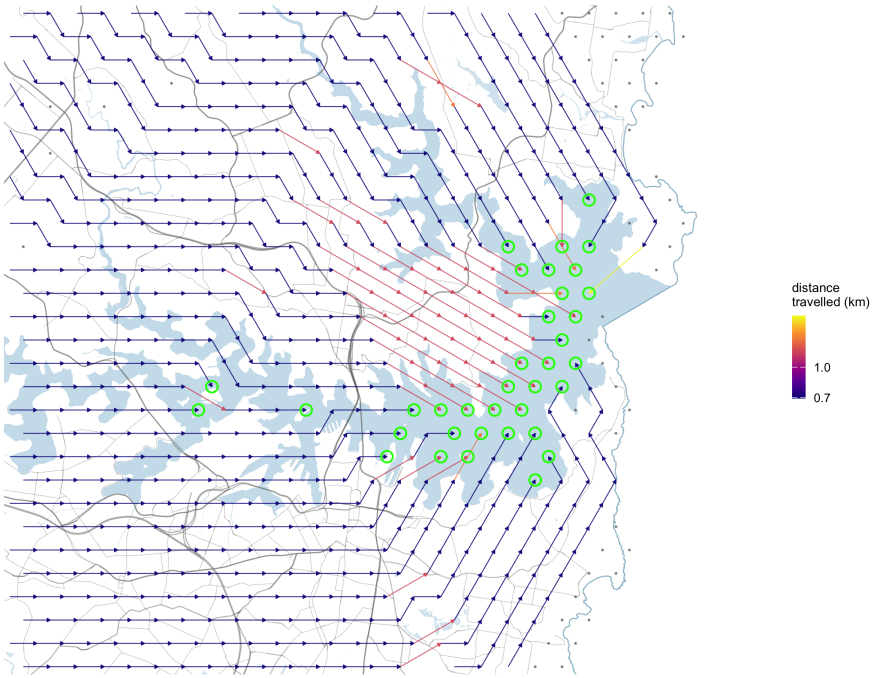


Figure 28: The paths taken by the hexagons in the harbour fill, zoomed to the estuary and its surrounds. Each arrow runs from a hexagon's origin site to the site the displacement-minimising bijection assigns it, coloured by distance travelled; green rings mark the 36 newly occupied harbour cells and grey dots hexagons that stay put. Because the cost is squared, a chain of short hops is cheaper than one long haul, so the harbour is filled by a conveyor of single-cell slides: of the 2,846 hexagons that move, 71% travel exactly one hexagon step (0.693 km) and none travels more than 10.8 km, even though the vacated cells lie 85–89 km away. Only 586 of those movers fall inside this frame — the cascade runs across the whole urban area (Figure 29).

In the baseline specification, the fundamental amenity, productivity and total built floorspace are shifted with the hexagon as they move– the city does not adjust to its new geography except by workers changing where they live and work (and accompanying price changes). To determine travel times in the new city, I first calculate travel times between hexes in and adjacent to the newly filled harbor using a travel time function of distance $f(d) = \alpha + \beta d$ estimated on existing inner city (but not harbor crossing) hexagon pairs. I then update the rest of the travel time matrix using a one step min-plus product of the access leg costs with the existing travel time matrix.

To summarize, the built form and underlying calibrated productivity and

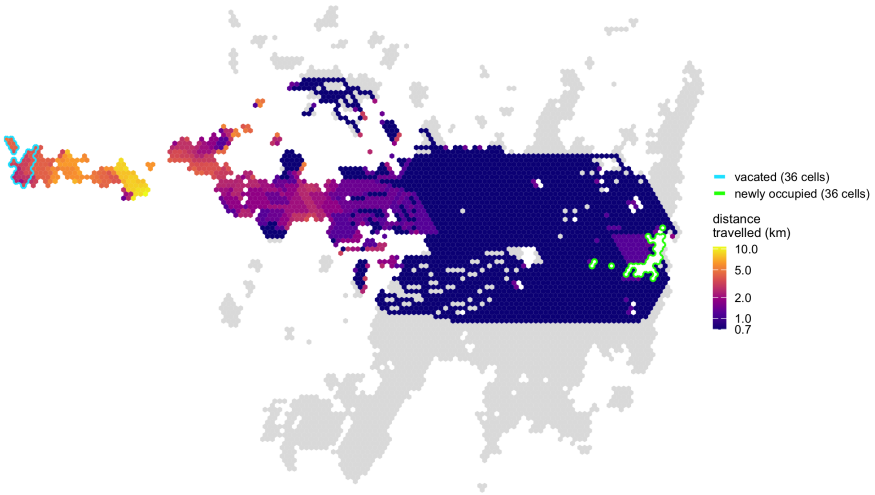


Figure 29: Distance travelled by each hexagon in the harbour fill, Sydney Significant Urban Area, drawn at each hexagon's origin site; grey cells stay put (42.7% of the 4,965 cells). The cyan outline rings the 36 vacated cells in the upper Blue Mountains and the green outline the 36 newly occupied cells in the harbour; the latter are target sites with no origin, so they are otherwise absent from the map. The reshuffle is a city-wide conveyor rather than a local rearrangement: a broad contiguous band of the city slides exactly one cell towards the harbour (dark blue), while the largest displacements are the local contraction of the far-western fringe around the 36 vacated cells, which are themselves moved only 3.4 km on average.

amenity of Sydney are shuffled conservatively inward until the harbor is filled, with the road network held fixed in the existing city and travel times inferred in the harbor.

A.3 Results

In the baseline specification where built floorspace is fixed and travels with the rest of the hexagon's attributes if it moves, welfare increases by 1.9% and commutes fall by 0.2% (from 20.96 to 20.91 minutes), while output *falls* by 0.6% (all compared to the baseline calibrated city)[22]. The welfare gain is driven by an increase in accessibility and amenity at the cost of production: filling the harbour in this way spreads employment out rather than concentrating it. The number of people in the highest employment 400m hexagon is essentially unchanged (up 0.5%, from 103,000 to 104,000), and that single hexagon is the one exception to the otherwise general pattern of dispersion: as shown in Figure 30 the employment share of the highest n hexagons is *lower* in the counterfactual at every rank beyond the first, the two cumulative curves crossing between the first hexagon and the second. The counterfactual's busiest hexagon holds 541 more workers than the baseline's, but its busiest two hold 4,551 fewer and its busiest hundred 33,000 fewer; employment in the densest $4\,\mathrm{km}^2$ falls by 7.7%. It therefore appears that geographic compaction of the city does not compact employment, because the locations whose accessibility improves most are not those that already enjoyed the best access and productivity. This is the opposite of the CBD shift of Section 4.1, which raised concentration at every rank. This may suggest one incidental benefit of the geographic barriers around the CBD: because there is less developable area arround these highly productive locations, firms are forced into a smaller area than they would otherwise be, and therefore increase agglomeration benefits to their neighbours.

Table 6: Harbour fill: welfare, output and commuting under the three specifications (Sydney SUA, 400 m hexagons, $\zeta = 1.86$). "Floorspace stock moves" (fixedH) is the specification described in the text; elastic lets floorspace supply respond; fixedB freezes residence amenities at their baseline values and is the amenity-endogeneity bracket. The two right-hand columns summarise the concentration reading discussed above.

	Δ welfare	Δ output	commute (min)		employment	densest
			baseline	c'factual	Gini	$4\,\mathrm{km}^2$
Floorspace stock moves, fixed (fixedH)	+1.87%	−0.60%	20.96	20.91	0.716	−7.7%
Floorspace supply elastic	+5.01%	−2.84%	20.96	20.94	0.703	−19.7%
Residence amenities frozen (fixedB)	+1.22%	−0.29%	20.96	20.89	0.728	−10.5%
Baseline city	—	—	20.96	—	0.724	—

[22]

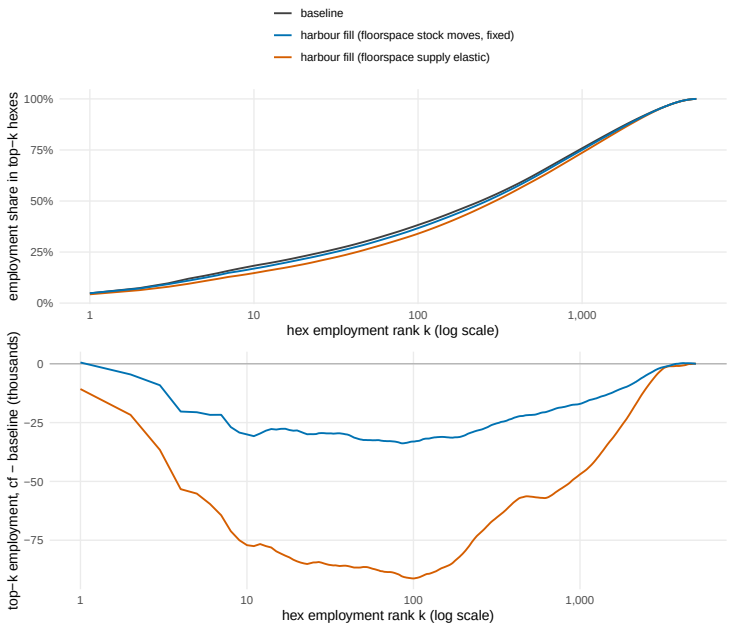


Figure 30: Employment concentration, baseline vs the harbour fill. Top panel: cumulative employment share of each city's top-k hexagons, with each scenario ranked by its own hexagon employment. Bottom panel: the difference between each counterfactual curve and the baseline. Unlike the CBD shift, the counterfactual curves lie *below* the baseline at essentially every rank and never cross back: filling the harbour makes the city less employment-concentrated throughout the distribution, with the largest shortfall 33,800 workers at $k = 86$ when the floorspace stock travels with the bundle, and 91,200 at $k = 100$ when floorspace supply is elastic.

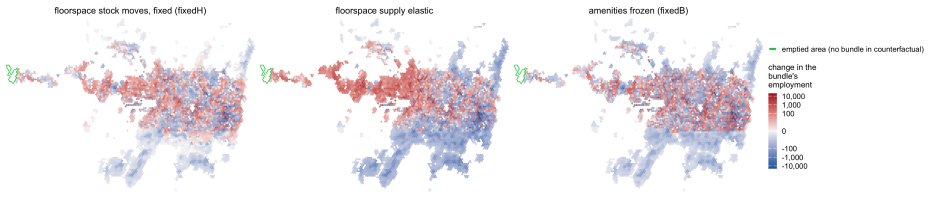


Figure 31: Employment change in the harbour fill, by bundle of attributes, under the three supply and amenity specifications. Each cell of the counterfactual footprint is coloured by the change in employment of the hexagon-bundle that occupies it there — its employment at its new site, minus the employment the same bundle had at its origin site in the baseline city — at the native 400 m resolution, on a shared signed-logarithmic diverging scale. The green perimeter rings the emptied area: the 36 vacated sites in the upper Blue Mountains, which hold no bundle in the counterfactual. The 36 bundles shuffled onto the newly occupied harbour cells (shown in detail in Figure 32) carried 30,200 workers at baseline and shed 16–36% of them in aggregate at their new sites, depending on the specification — a loss concentrated almost entirely in the single largest arriving bundle, while most of the other arrivals gain; across all bundles the map shows a broad transfer of employment from the city's east towards its middle and west.

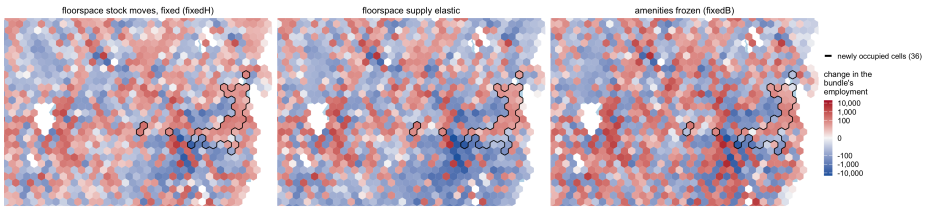


Figure 32: Zoom of Figure 31 to the estuary (the estuary's bounding box plus 2 km; the emptied area of the full map lies far outside this window). The black outline rings the 36 newly occupied cells; the remaining light-blue water is the part of the estuary the fill leaves open. Most cells inside the fill read red: 30 of the 36 arriving bundles gain employment, the median being a foreshore cell of roughly a hundred workers that adds several dozen. The fill nevertheless sheds employment in aggregate — 16% of its 30,200 baseline workers when the floorspace stock moves with them, 32% with amenities frozen, 36% under elastic supply — because two thirds of that baseline employment travels in a single bundle: the 20,500-worker Circular Quay cell at the CBD's harbour front, nudged one cell eastwards into the harbour, loses 30–53% of its workers at its new site. The reclaimed harbour thus keeps, but shrinks, the CBD's foreshore apex, and does not become a new employment cluster in any specification.

A.4 Compacting to a circle

Filling the harbor reduces the distortion of Sydney's shape by increasing space in the most central areas. To take the excercise to its logical endpoint, I here reshuffle the city to fill a perfect circle. This is implemented analogously to the harbor fill counterfactual: I set a new target footprint whose center is chosen to maximise its overlap with the existing currently occupied set. This center is 4.3km South East of Parramatta, 16.1 km west of the current CBD, and overlaps 3,504 of the existing 4,965 sites (70 percent) (whereas a disc centered at the current cbd would only overlap 2,599 sites, 52 percent, with over 42.5% of the disc falling in open water). I then solve for the set of permutations of hexagons that minimimises the sum of squared displacements to find the least disruptive change, and impute travel times through previously occupied land using the same function of distance as before.

Results: in the specification where the built floorspace stock travels with the hexagon, output falls by 2.1%, average commutes rise by 2.2% (from 20.96 to 21.43 minutes), and welfare increases by 5.6%. Employment deconcentrates sharply: the highest employment hexagon falls by 12% (from 103,000 to 91,000 workers) and employment in the densest 4 km^2 falls by 21%. Allowing floorspace supply to re-equilibrate at each hexagon's new location raises the welfare gain to 22.6%, with commutes rising by 26% and output falling by 11.8%; holding both floorspace supply and residence amenities fixed at their baseline values attenuates welfare gain to 5.6%, and output drop to 2.848%. Figure 33 shows the reshuffle: 99.8% of hexagons move, by a median of 4.5 km and a maximum of 48 km.

A.5 part superseded: New transport lines

How much of the costs of its shape can be attributed to the extra cost of transport across the bay? I analyse this by considering i) the effect of completing the new west harbor tunnel, ii) the effect of removing all major harbor-adjacent public transport improvements of the last decade

Existing Results for this are currently based on old model specifications. The effects were very small, but this may be due to the fact that the current model does not feature congestion douting. I'll Probably leave a rerun of this until I have congestion routing in place.

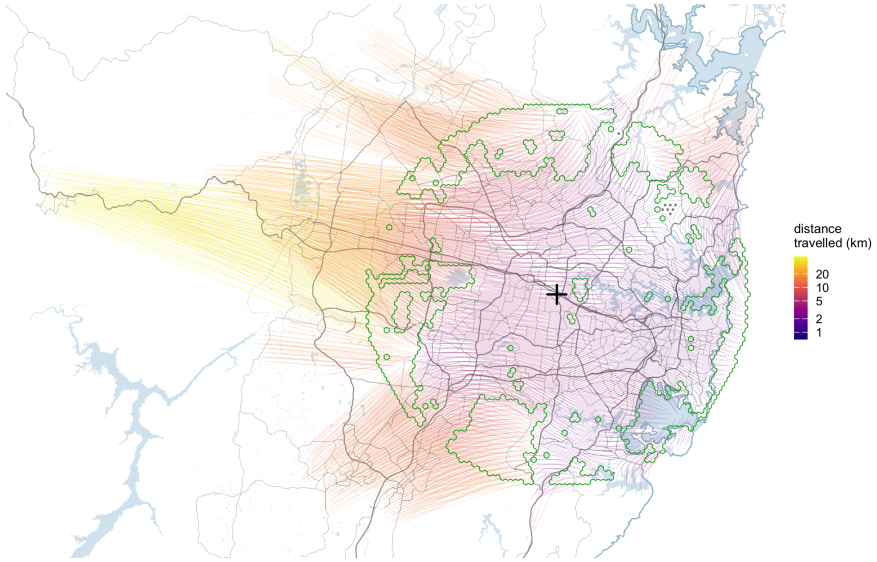


Figure 33: The circle-city reshuffle. Each line runs from a hexagon's origin site to the site assigned by the displacement-minimising bijection, coloured by distance travelled. The green line traces the perimeter of the newly occupied area — the 1,461 positions the compact city takes over that lie outside today's footprint, which is why it is not a simple circle: it follows the disc's outer edge only where the disc reaches beyond the existing city, and runs internal boundaries wherever the city already covered it. The cross marks the centre of the target disc. Almost the whole city moves (99.8% of hexagons, median 4.5 km), with the far-western corridor drawn 20–48 km eastwards. The overlap-maximising disc has a radius of 25.7 km and is centred 4.3 km south-east of Parramatta.

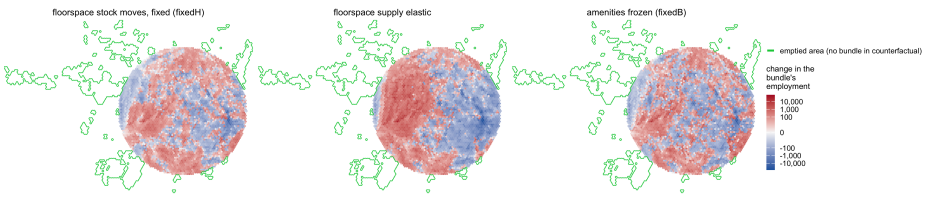


Figure 34: Employment change in the circle-city reshuffle, by bundle of attributes, under the three supply and amenity specifications. Construction as in Figure 31: each cell of the target disc is coloured by the change in employment of the hexagon-bundle that occupies it there — counterfactual employment at its new site minus baseline employment at its origin site — at the native 400 m resolution on a shared signed-logarithmic diverging scale. The green perimeters ring the emptied areas: the 1,461 sites of today's footprint that lie outside the disc and hold no bundle in the compact city. Read this way the disc is not a wash of gains: the 1,461 bundles relocated onto previously unoccupied ground — carrying 338,000 baseline workers, five sixths of them on the disc's western side — gain 19–21% in aggregate in every specification, while the deepest losses are the old CBD's largest bundles, which move only a few kilometres and remain in the city's east but shed up to 14,500 workers each when the floorspace stock moves with them and up to 73,100 under elastic supply.

B Additional employment exhibits

Employment series for the Parramatta circle and its comparators. The government timeline and the private-response pair, which the results section cites, come first (Figures 35–37); the series after them are not referenced directly. The public/private split is in the main text as Figure 20.

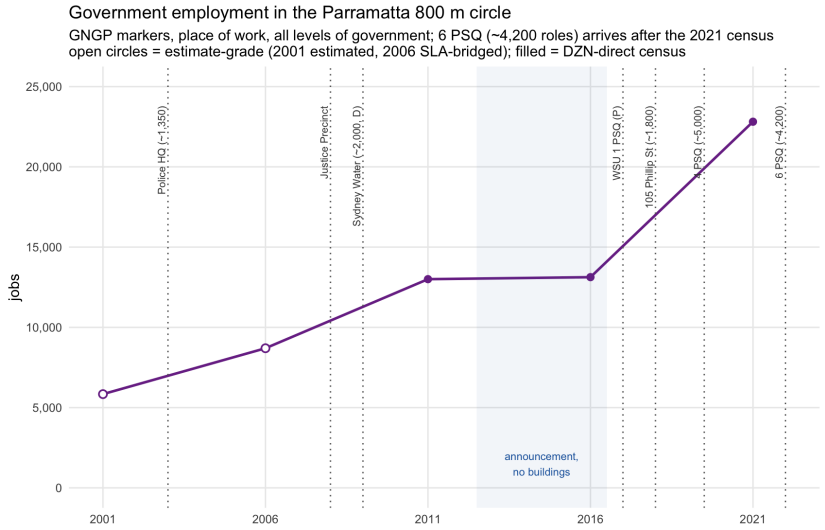


Figure 35: Government employment in the Parramatta 800 m circle against the documented office moves (dotted verticals). The series is the census GNGP government-employer markers at place of work, all levels of government combined (2001–2021); the movement is almost entirely NSW state government, the level the relocation policy moves. Open markers are estimate-grade years (2001 estimated from the industry mix, 2006 bridged from SLA-level data); filled markers are direct Destination-Zone measurements. Government employment is flat between the 2011 and 2016 censuses – the announcement-without-buildings window – then rises by 9,684 (state: 8,548) to 2021, a rise the census still half-censors: 6 PSQ's ∼4,200 roles arrived after census day and TZP24 places the census-2021 level ∼19% below its recalibrated estimate.

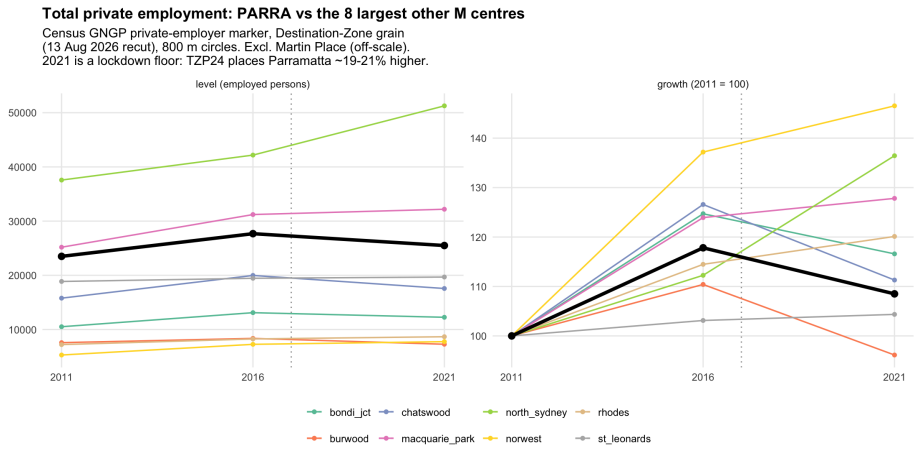


Figure 36: Total private employment: Parramatta against the eight largest other Division-M centres (800 m circles). Left: levels; right: indexed to 2011 = 100. The series is the census GNGP private-employer marker at place of work, which exists at Destination-Zone grain only for 2011/2016/2021 – hence three points. Parramatta (bold) is mid-pack: $\times 1.085$ over 2011–2021, 13th of the 24 units in pool V1. The Sydney CBD (Martin Place) is excluded as off-scale, and the plotted comparators are selected on Division-M size, so they are not a neutral sample – the ranks quoted in the text are computed on the full V1 pool, not on these nine lines. The 2021 point is a lockdown floor, as in Figure 35.

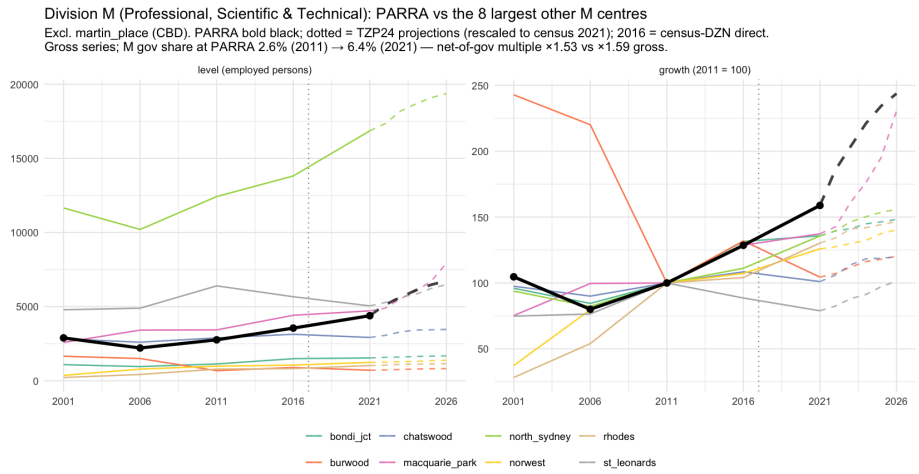


Figure 37: ANZSIC Division M (professional, scientific and technical services) employment: the same nine centres and the same colours as Figure 36, over the full 2001–2026 span the division-level series supports. Parramatta grows ×1.59 over 2011–2021, first of the 24 units in pool V1; net of government employers within M the multiple is ×1.53 and the rank third. Dotted segments from 2022 are TZP24 projections rescaled to census 2021. Read against Figure 36, the contrast is the finding: mid-pack in private employment overall, first in professional and technical services.

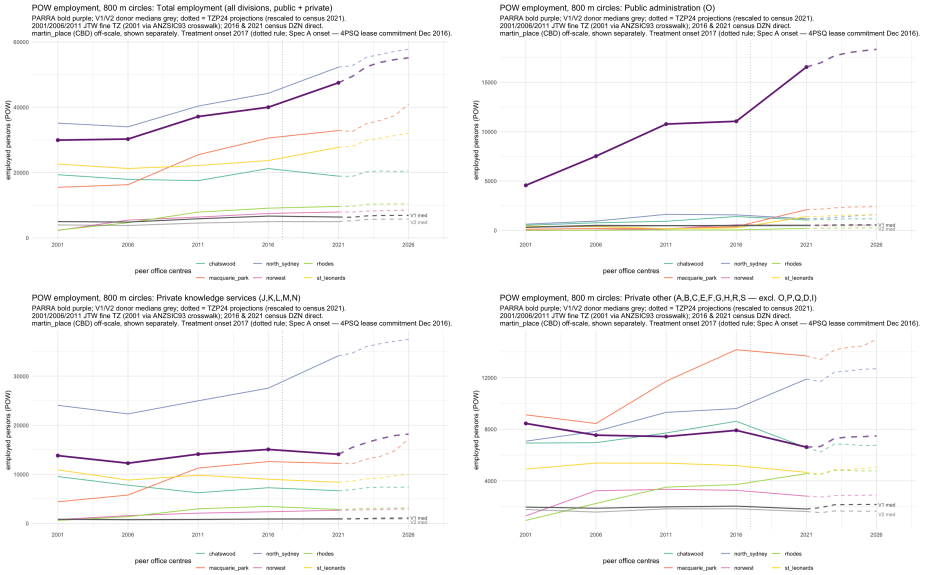


Figure 38: Employment in the Parramatta circle vs comparator subcentres (800 m circles, fine-geometry attribution). Top: total and public administration (ANZSIC Division O). Bottom: the private-sector proxy buckets – knowledge services (J–N) and other private (A–C, E–H, R, S). Education & health (P, Q) and utilities & transport (D, I) are excluded from the private buckets because they are heavily government at Parramatta.

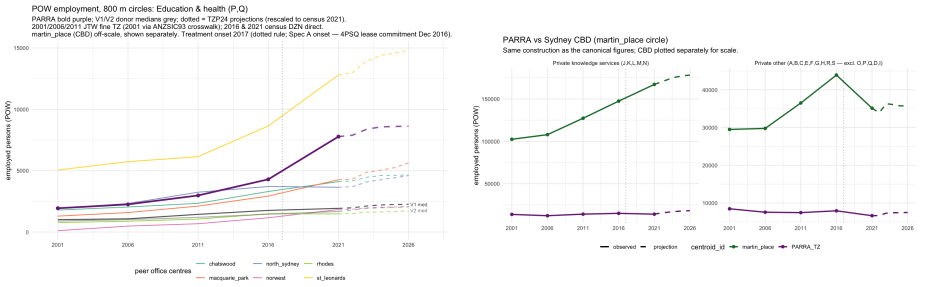


Figure 39: Left: education & health employment (Divisions P, Q) – the WSU vertical-campus arrival (2017) is the visible break. Right: Parramatta vs the Sydney CBD (Martin Place circle, own scale).

C Dynamic Model with a public sector

To understand the dynamic equilibrium response of the city to the goverment office moves, I extend my static model to include dynamic worker and construction decisions, and a public sector. Two key facts from the reduced form govern our approach: valuations in parrmatta appear to substantially increase compared to control units from 2017 onward, and private employment in parramatta appears relatively flat compared to controls beteween 2016-2021. To capture the fact that workers' locations may be sticky over time, I follow Takeda and Yamagishi 2026 by introducing forward looking workers with Calvo shock migration opportunities, with the frequency of moves calibrated to census data. Within periods, workers' place of work and residence is fixed, and they receive flow utility from the amenity of their place of residence less rental costs and the wages in their place of work less commute costs. Between periods, workers who are granted a migration opportunity choose where to live and work to maximise expected lifetime utility, subject to a gumbel preference shock. To capture the fact that valuations can capitalise future shocks on announcement, I augment the model of Takeda and Yamagishi 2026 by having landlowners respond to the full path of future rents. Floorspace decays each period, and new construction follows as a constant elasticity function of expected lifetime rental returns. Finally, I add to Takeda and Yamagishi 2026 in two further ways to capture features relevant for my policy analysis: multiple worker types (as in my static model), and a public sector.

The public sector is seperate from the private sectors, and controlled by a policy maker that dictates where to rent offices for its workers. The key parameter determining the effects of the government move is the local spillover from public to private sector firms: if these spillovers are high, the government can alter the distribution of private sector jobs across the city just by moving its own workers. If the spillovers are higher for low income jobs, or the effects of shifting the CBD West are more positive for low income workers, then they can reduce inequality by moving public sector jobs West.

The result is a model in which floorspace values respond immediately to increases in future lifetime profits, whereas population distributions take time to adjust. The model is tractable and fast to solve because the forward looking component of worker decisions is unnaffected by their prior locations, and capital decisions are simple functions of local future rents. The long run equilibrium is identical to the static model described earlier with an added public sector. The dynamic paramaters, and hence relative speeds of response, are calibrated Sydney wide moments and values from the literature, and the office move shock is used to calibrate the public-private spillover. At present, I calibrate the model to 2016 data, run it forward under the presence of a public sector move to Parramatta, and search for the value of spillovers that leads to short run model responses matching the observed reduced form treatment effects. I am in the process of adding a full dynamic estimation of the model.

C.1 Workers

C.1.1 Within period

Within period, workers' location of residence and workplace is fixed. Their utility is a cobb douglas function of floorspace, goods consumption and local amenities, which leads to the following log indirect flow utility, the same as the static model:

$$\ln u_{ijt}^g = \ln B_{it}^g - (1-\beta^g)\ln Q_{it} + \ln w_{jt}^g - \ln d_{ij}^g \tag{38}$$

Given a fixed within period number of workers living in i and commuting to j L_{ij}^g, total worker expenditure on floorspace in i is:

$$Q_i H_{Ri} = \sum_{j,g} (1-\beta^g) w_j^g L_{ij}^g \tag{39}$$

C.1.2 Between period

Between periods, workers are granted the opportunity to move with probability θ^g. Workers that are allowed to move receive a Gumbel residence-workplace location pair preference shock and choose their location of residence and workplace to maximise their lifetime utility with rational expectations (i.e. correctly anticipating the full path of future prices).

Lifetime utility is given by:

$$\mathcal{V}_{ijt}^g = \underbrace{\ln u_{ijt}^g}_{\text{flow utility}} + \underbrace{(1-\theta^g)\,\iota \mathcal{V}_{ijt+1}^g}_{\substack{\text{continuation value} \\ \text{if stuck}}} + \underbrace{\theta^g E\big[\max_{i'j'} \iota \mathcal{V}_{i'j't+1}^g + z_{i'j't+1}\big]}_{\substack{\text{continuation value} \\ \text{of movers}}} \tag{40}$$

Where

- $z_{i'j't+1}$ is the idiosyncratic preference shock upon being granted the opportunity to move, drawn on the discounted continuation value, and is Gumbel distributed with scale parameter σ^g.
- ι is the discount on the value of future periods

Given the properties of the Gumbel distribution, the expecation value can be written as[23]:

$$\mathcal{V}_{ijt}^g = \ln u_{ijt}^g + \iota(1-\theta^g)\mathcal{V}_{ijt+1}^g + \theta^g \sigma^g \ln\Big(\sum_{i'j'} e^{\iota \mathcal{V}_{i'j't+1}^g / \sigma^g}\Big) \tag{41}$$

[23](dropping the additive constant $\theta^g \sigma^g \gamma_E$, with γ_E the Euler–Mascheroni constant, which is common to all pairs)

The share of those with the opportunity to move in $t+1$ who choose to live in i and work in j, denoted m^g_{ijt+1}, has the Gumbel share form

$$m^g_{ijt+1} = \frac{e^{\iota \mathcal{V}^g_{ijt+1}/\sigma^g}}{\sum_{i'j'} e^{\iota \mathcal{V}^g_{i'j't+1}/\sigma^g}} \tag{42}$$

Net entrants to the city are assumed to choose a residence–workplace pair in the same way as the incumbent movers. With net entrants given by the expression $L^g_{\text{tot},t+1} - L^g_{\text{tot},t}$, the law of motion is therefore:

$$L^g_{ijt+1} = (1-\theta^g) L^g_{ijt} + \Big[\theta^g L^g_{\text{tot},t} + \big(L^g_{\text{tot},t+1} - L^g_{\text{tot},t}\big)\Big] m^g_{ijt+1} \tag{43}$$

Value function simplification: The value function is nominally an $N \times N \times T$ object, one number for every residence–workplace pair in every period. However, the fact that movers' residence-location choice is unnaffected by their prior location means the value function actually inherits the residence-workplace separability of flow utiltity. The continuation value term for movers in (41) — the log-sum over all pairs — is the same for every ij, because a worker granted a moving opportunity chooses freely over the whole set of pairs, with no cost that depends on where they came from:

$$\Omega^g_t \equiv \theta^g \sigma^g \ln\Big(\sum_{i'j'} e^{\iota \mathcal{V}^g_{i'j't+1}/\sigma^g}\Big), \tag{44}$$

Meanwhile, if we expand out the future sum of stuck continuation values in terms of the sum of each periods flow utility [24], we see that every term inside that sum is a function of the residence i alone, of the workplace j alone, or of the exogenously fixed commute. Thus if we define $\mathcal{R}^g_{it}$ as the present value of the residential side of flow utility at i and $\mathcal{F}^g_{jt}$ as the present value of log wages at j:

$$\begin{aligned} \mathcal{R}^g_{it} &\equiv \sum_{s\geq 0} \big[\iota(1-\theta^g)\big]^s \big[\ln B^g_{it+s} - (1-\beta^g)\ln Q_{it+s}\big], \\ \mathcal{F}^g_{jt} &\equiv \sum_{s\geq 0} \big[\iota(1-\theta^g)\big]^s \ln w^g_{jt+s}, \end{aligned} \tag{45}$$

We see that the value function can be rewritten as:

$$\mathcal{V}^g_{ijt} = \mathcal{R}^g_{it} + \mathcal{F}^g_{jt} - \frac{\kappa^g \tau_{ij}}{1-\iota(1-\theta^g)} + \bar{\Omega}^g_t. \tag{46}$$

The $N \times N \times T$ array is therefore generated by two $N \times T$ series, the single fixed $N \times N$ commute cost matrix, and the $1 \times T$ option value. Thus the unknowns of the forward problem are $2N$ numbers per period per type — the anticipated paths of Q_{it} and w^g_{jt} — and not the N^2 that a general $\mathcal{V}^g_{ijt}$ would require.

[24] possible since $\iota(1-\theta^g) < 1$ and flow utility is bounded

In reality, Individuals sometimes move either their job or their residence one at a time rather than moving both. The above setup with a single shock probability leading to choice of both simultaneously will both lead to an over/underestimate of the movement frequency of either residence or workplace, and miss the fact that eg the choice of new workplace is often constrained by current residence. To capture these facts, In an extension in the appendix I provide a model with separate Calvo shocks for each of workplace and residence individually and the pair jointly.

C.2 Production

Firms operate myopically each period solving a static problem identical to that described in the static model section, so the key equations are the same: C5pommercial floorspace demand:

$$H_{Fj}^t = L_{Fj}^t \Big((1-\alpha) \frac{A_j^t}{Q_j} \Big)^{1/\alpha} \tag{47}$$

Type t wage offers at j as a function of local productivity and rents:

$$w_j^t = \alpha (1-\alpha)^{\frac{1-\alpha}{\alpha}} (A_j^t)^{\frac{1}{\alpha}} Q_j^{-\frac{1-\alpha}{\alpha}} \tag{48}$$

C.3 Public Sector

Public sector workers come in the same two education types as private sector workers. Each period, the policymaker decides how many public sector employees of each type to hire $L_{\text{Ptot},t}^g$, where they will work L_{PFit}^g, and their wage w_{Pt}^g (type-specific, but fixed across all locations in this basic version). Where no type superscript appears below, public sector objects denote type sums: $L_{PFjt} \equiv L_{PFjt}^h + L_{PFjt}^l$, $L_{PRit} \equiv L_{PRit}^h + L_{PRit}^l$, $L_{\text{Ptot},t} \equiv L_{\text{Ptot},t}^h + L_{\text{Ptot},t}^l$. The government then hires floorspace in each location such that the expenditure share on floorspace matches the private sector firm expenditure share on floorspace:

$$H_{PFjt} = \frac{1-\alpha}{\alpha} \sum_g w_{Pt}^g \, L_{PFjt}^g / Q_{jt} \tag{49}$$

In the current simple version of the model, public sector workers have fixed location of residence, location of workplace and wage (which I calibrate in the final period). In their place of residence they demand floor-space such that their expenditure share matches that of private sector workers of the same education type:

$$H_{PRit} = \sum_g (1-\beta^g) \, w_{Pt}^g \, L_{PRit}^g / Q_{it} \tag{50}$$

The outlays on the public sector — the wage bill, its office rents, and any rent subsidies S_{jt} — are funded from outside the city in this basic version.

C.4 Floorspace Supply

I assume that floorspace is perfectly substitutable between commercial and residential uses and has a common price Q_{it} before any taxes/subsidies given to producers. Within period, floorspace supply is fixed. Rents adjust so that floorspace market clearing is satisfied:

$$\begin{aligned} H_{it} &= H_{Rit} + H_{Fit} + H_{PRit} + H_{PFit} \\ &= \sum_g \Big[\sum_j \frac{(1-\beta^g)}{Q_{it}} w_{jt}^g L_{ijt}^g + L_{Fit}^g \Big((1-\alpha) \frac{A_{it}^g}{(1-S_{it})Q_{it}} \Big)^{1/\alpha} + \frac{(1-\beta^g)\, w_{Pt}^g\, L_{PRit}^g}{Q_{it}} + \frac{1-\alpha}{\alpha} \frac{w_{Pt}^g\, L_{PFit}^g}{Q_{it}} \Big] \end{aligned} \tag{51}$$

where $L_{Fit}^g = \sum_r L_{rit}^g$ is type-g private employment at i, predetermined within the period by the commute-pair distribution, and the four terms are private residential demand, private commercial demand, public residential demand and public office demand respectively.

Between periods, floorspace decays and forward looking absentee landlords decide how much new floorspace J_{it} to build.

$$f_{it+1} = (1-\delta_f) f_{it} + J_{it}$$

To construct new floorspace, landlords face a constant elasticity cost function. The cost increases as a function of floorspace built this period, rather than a function of the total existing floorspace supply, representing the fact that resources and opportunities to build each period are limited. As seen below, this ensures a simple closed form expression for expenditure on new floorspace. It results in an elasticity of supply that is lower in the short run and higher in the long run. The cost is given by

$$c_i(J_{it}) = \frac{\zeta}{1+\zeta} \psi_i^{-1/\zeta} J_{it}^{1+1/\zeta} \tag{52}$$

where:

- J_{it} is floorspace constructed per unit land in period t
- ζ is the elasticity of supply
- ψ_i is a location specific cost shifter, which could include costs due to zoning regulation or topography

Because developers respond to lifetime returns from new floorspace (rather than only next periods rent as in Takeda and Yamagashi), development will respond to future shocks (such as the announcement of future office moves), and the model produces expressions for land value that I can compare to my land value data.

A unit of floorspace built in t is first let in $t+1$ and then decays at rate δ_f. Landlords discount future rents at ι_f, which plays the same role as ι in

the worker's value function, so the present value of the discounted future profits from constructing J_{it} in t is:

$$\pi_{it} = J_{it} \sum_{s=t+1}^{\infty} \iota_f^{s-t} (1-\delta_f)^{s-t-1} Q_{is} - \frac{\zeta}{1+\zeta} \psi_i^{-1/\zeta} J_{it}^{1+1/\zeta} \tag{53}$$

If we define the present value of the lifetime profits from a unit of constructed floorspace as:

$$P_{it} = \sum_{s=t+1}^{\infty} \iota_f^{s-t} (1-\delta_f)^{s-t-1} Q_{is} \tag{54}$$

then the floorspace FOC gives:

$$J_{it} = \psi_i P_{it}^{\zeta} \tag{55}$$

Therefore we have that:

$$f_{it+1} = (1-\delta_f) f_{it} + \psi_i P_{it}^{\zeta}$$

note that to find the new total floorspace constructed we need to multiply this by the available land K_i

$$H_{it+1} = (1-\delta_f) H_{it} + K_i \psi_i P_{it}^{\zeta} \tag{56}$$

The value per unit area of currently empty (unimproved) land UV^{land} is the value of current and future profits from construction on that land, discounted at the landlords' ι_f (the sum starts at $s = t$ so that, like the rent term below, the value is cum-dividend: it includes the profit from construction in the current period):

$$UV_{it}^{land} = \sum_{s \geq t} \iota_f^{s-t} [J_{is} P_{is} - \frac{\zeta}{1+\zeta} \psi_i^{-1/\zeta} J_{is}^{1+1/\zeta}]$$

which with the above profit maximising FOC gives:

$$UV_{it}^{land} = \frac{\psi_i}{1+\zeta} \sum_{s \geq t} \iota_f^{s-t} P_{is}^{1+\zeta} \tag{57}$$

The total value of land with existing floorspace is the sum of the value of the existing floorspace plus the option value from future construction. As the amount constructed each period is not dependent on existing construction (except indirectly via Q), this means we can sum the value of existing floorspace and the above expression for unimproved land value

$$V_{it}^{land} = f_{it}\big(Q_{it} + (1-\delta_f) P_{it}\big) + \frac{\psi_i}{1+\zeta} \sum_{s \geq t} \iota_f^{s-t} P_{is}^{1+\zeta} \tag{58}$$

where the first term is the current rent on the existing stock plus the value of the $(1-\delta_f) f_{it}$ units that survive into $t+1$, each worth P_{it}.

C.5 Agglomeration forces

Agglomeration of workers causes local productivity spillovers. Each receiving type may benefit differently from agglomeration, and each source group — low-type private workers and public sector workers, relative to high-type private workers — may contribute differently to it. Productivity spillovers depend on a single weighted employment access term,

$$A_{jt}^g = a_j^g \left(\Upsilon_{jt}^A\right)^{\lambda_g}, \qquad \Upsilon_{jt}^A = \sum_s e^{-\delta \tau_{js}} \frac{L_{Fst}^h + \phi_l L_{Fst}^l + \phi_P L_{PFst}}{K_s} \tag{59}$$

and amenity spillovers on a weighted residence access term,

$$B_{it}^g = b_i^g \left(\Upsilon_{it}^B\right)^{\eta_g}, \qquad \Upsilon_{it}^B = \sum_s e^{-\rho \tau_{is}} \frac{L_{Rst}^h + \phi_l^B L_{Rst}^l + \phi_P^B L_{PRst}}{K_s} \tag{60}$$

The source weights ϕ_l, ϕ_P (and ϕ_l^B, ϕ_P^B on the residence side) govern how much each group contributes to the spillover relative to high-type private workers, whose weight is normalised to one; the type-specific exponents λ_g, η_g govern how strongly each receiving type benefits. I choose a simple sum of public and private workers rather than product, because a product would lead to large gains from the first few public sector workers. In future I may test robustness to a CES function.

The location-and-type fundamentals a_j^g and b_i^g are time-invariant; the composites A_{jt}^g, B_{it}^g are endogenous and vary over time, because the access terms depend on the period-t equilibrium employment and residence distributions. This is a second channel coupling the two types, in addition to the floorspace market: a location's employment of any source group raises both firm types' productivity there (all sources enter the one access term), and likewise for residence and amenities.

C.6 Equilibrium

C.6.1 Within period

Each period, there is a static equilibrium of the model taking the number of workers in each residence-workplace location pair L_{ij}^g and floorspace supply H_i as given.

Given parameters

$$\Theta = \{\alpha, \beta^g, \lambda_g, \delta, \eta_g, \rho, \phi_l, \phi_P, \phi_l^B, \phi_P^B\},$$

and exogenous fundamentals

$$\{a_i^g, b_i^g, \psi_i, K_i, \tau_{ij}, S_i\},$$

and the state variables L_{ijt}^g and H_{it}, the within period static equilibrium is the vector

$$\{Q_i, w_i^g, A_i^g, B_i^g\}$$

which satisfies the following equations:

The composite productivities and amenities A_i^g, B_i^g follow immediately from the state variable resident and employment populations via agglomeration equations 59–60.

Wages are a simple closed form function of floorspace rents Q_{jt} given profit maximisation and zero profits:

$$w_{jt}^g(Q_{jt}) = \Lambda \left(A_{jt}^g\right)^{1/\alpha} \big((1 - S_{jt})Q_{jt}\big)^{-(1-\alpha)/\alpha}, \qquad \Lambda \equiv \alpha(1-\alpha)^{(1-\alpha)/\alpha} \tag{61}$$

Floorspace rents solve the floorspace market clearing condition $Q_{it}H_{it} = X_{it}(Q)$ with

$$\begin{aligned} X_{it}(Q) = (1 - t_{w,t}) \Big[\sum_g \sum_j (1 - \beta^g)\, w_{jt}^g(Q_{jt})\, L_{ijt}^g \\ + \sum_g (1 - \beta^g)\, w_{Pt}^g L_{PRit}^g \Big] \\ + \sum_g \frac{1-\alpha}{\alpha} \frac{w_{it}^g(Q_{it})\, L_{Fit}^g}{1 - S_{it}} + \frac{1-\alpha}{\alpha} \sum_g w_{Pt}^g L_{PFit}^g. \end{aligned} \tag{62}$$

In short, once we find the floorspace price that solves the contraction mapping of floorspace market clearing, the other within period variables follow straightforwardly by closed form expressions.

C.6.2 Dynamic equilibrium

Given the parameters

$$\Theta = \{\alpha, \beta^g, \kappa^g, \zeta, \lambda_g, \delta, \eta_g, \rho, \phi_l, \phi_P, \phi_l^B, \phi_P^B, \iota, \iota_f, \delta_f, \theta^g, \sigma^g\},$$

and exogenous fundamentals

$$\{a_i^g, b_i^g, \psi_i, K_i, \tau_{ij}, S_i\},$$

and an exogenous paths[25] for all t of variables $L_{tot,t}^g, L_{PFjt}^g, L_{PRit}^g, w_{Pt}^g, S_{jt}, t_{w,t}$, and initial conditions L_{ij0}^g, H_{i0},

the dynamic equilibrium is the sequence of variables:

$$L_{ijt}^g, \mathcal{V}_{ijt}^g, H_{it}, P_{it}$$

that solve the following dynamic equations:

Worker Bellman equation:

$$\mathcal{V}_{ijt}^g = \ln B_{it}^g + \ln w_{jt}^g - (1-\beta^g) \ln Q_{it} - \ln d_{ij}^g + \iota(1-\theta^g)\mathcal{V}_{ijt+1}^g + \theta^g \sigma^g \ln\Big(\sum_{i'j'} e^{\iota \mathcal{V}_{i'j't+1}^g / \sigma^g}\Big) \tag{63}$$

[25] restricted so that the sums defining $\mathcal{V}_t and P_t$ converge

Population evolution:

$$L^g_{ijt+1} = (1-\theta^g)L^g_{ijt} + \Big[\theta^g L^g_{\text{tot},t} + \big(L^g_{\text{tot},t+1} - L^g_{\text{tot},t}\big)\Big] \frac{e^{\iota \mathcal{V}^g_{ijt+1}/\sigma^g}}{\sum_{i'j'} e^{\iota \mathcal{V}^g_{i'j't+1}/\sigma^g}} \quad (64)$$

Floorspace supply evolution in each location:

$$H_{it+1} = (1-\delta_f)H_{it} + K_i \psi_i P_{it}^{\zeta}$$

where the Developers floorspace value is:

$$P_{it} = \sum_{s=t+1}^{\infty} \iota_f^{s-t} (1-\delta_f)^{\,s-t-1} Q_{is} \quad (65)$$

along with the set of static variables that solve the within period static equilibria above.

C.7 Steady state

In a steady state the fundamentals, the public sector vectors $L^g_{PFj}, L^g_{PRi}, w^g_P$, the policy instruments and the total population of each type L^g_{tot} (at its terminal value) are constant, and all endogenous objects are time-invariant. Because the option value of a future move is common to every pair, differencing the Bellman equation 41 across pairs and dropping time subscripts gives

$$\mathcal{V}^g_{ij} = \frac{\ln u^g_{ij} + \theta^g \sigma^g \ln\big(\sum_{i'j'} e^{\iota \mathcal{V}^g_{i'j'}/\sigma^g}\big)}{1-\iota(1-\theta^g)}, \qquad \mathcal{V}^g_{ij} - \mathcal{V}^g_{kl} = \frac{\ln u^g_{ij} - \ln u^g_{kl}}{1-\iota(1-\theta^g)} \quad (66)$$

so that lifetime value differences are flow utility differences scaled by the expected tenure of a pair. Substituting into the choosers' share 42, the common option value cancels and

$$m^g_{ij} = \frac{(u^g_{ij})^{\varepsilon^g_{\text{SS}}}}{\sum_{i'j'} (u^g_{i'j'})^{\varepsilon^g_{\text{SS}}}}, \qquad \varepsilon^g_{\text{SS}} \equiv \frac{\iota}{\sigma^g\big[1-\iota(1-\theta^g)\big]} \quad (67)$$

The stationary commute-pair distribution solves $L^g_{ij} = (1-\theta^g)L^g_{ij} + \theta^g L^g_{\text{tot}} m^g_{ij}$, i.e.

$$L^g_{ij} = L^g_{\text{tot}}\, m^g_{ij} = L^g_{\text{tot}} \frac{\big(B^g_i\, w^g_j\, Q_i^{-(1-\beta^g)}\, e^{-\kappa^g \tau_{ij}}\big)^{\varepsilon^g_{\text{SS}}}}{\sum_{r,s}\big(B^g_r\, w^g_s\, Q_r^{-(1-\beta^g)}\, e^{-\kappa^g \tau_{rs}}\big)^{\varepsilon^g_{\text{SS}}}}, \qquad L^g_{Ri} = \sum_j L^g_{ij}, \quad L^g_{Fj} = \sum_i L^g_{ij} \quad (68)$$

which is the commuting matrix of the static Ahlfeldt et al model with the Fréchet shape parameter replaced by $\varepsilon^g_{\text{SS}}$ (the residence and workplace scale parameters T_i, E_j of that model correspond here to pair-separable location shifts in the Gumbel distribution, and are absorbed into B^g_i and w^g_j exactly as in the adjusted variables of the calibration).

In steady state the present value of the lifetime rents earned by a unit of floorspace is:

$$P_i = \frac{\iota_f \, Q_i}{1 - \iota_f(1-\delta_f)} \tag{69}$$

which along with the profit maximising construction choice $J_{it} = \psi_i P_{it}^{\zeta}$ and stationarity of the stock, $\delta_f H_i = K_i J_i$, gives

$$H_i = \frac{K_i \, \psi_i}{\delta_f} \, P_i^{\zeta} = K_i \, \tilde{\psi}_i \, Q_i^{\zeta}, \qquad \tilde{\psi}_i \equiv \frac{\psi_i}{\delta_f} \Big(\frac{\iota_f}{1 - \iota_f(1-\delta_f)} \Big)^{\zeta} \tag{70}$$

We also have that the steady state unimproved land value per unit area is:

$$UV_i^{land} = \frac{\psi_i}{1+\zeta} \sum_{s \geq t} \iota_f^{s-t} P_i^{1+\zeta} = \frac{\psi_i}{(1+\zeta)(1-\iota_f)} P_i^{1+\zeta} = \frac{\delta_f \, \iota_f \, \tilde{\psi}_i \, Q_i^{1+\zeta}}{(1+\zeta)(1-\iota_f)\big(1 - \iota_f(1-\delta_f)\big)} \tag{71}$$

where the last equality substitutes the steady state price 69 and the calibrated $\tilde{\psi}_i$ of 70. Using the stationary stock $f_i = \psi_i P_i^{\zeta}/\delta_f$ and the steady state price, which implies $Q_i + (1-\delta_f)P_i = P_i/\iota_f$ (the cum-dividend value of an existing unit is the developer's value shifted one period earlier), the steady state total land value per unit area is:

$$V_i^{land} = f_i\big(Q_i + (1-\delta_f)P_i\big) + UV_i^{land} = \psi_i P_i^{1+\zeta} \Big[\frac{1}{\delta_f \, \iota_f} + \frac{1}{(1+\zeta)(1-\iota_f)} \Big] = \frac{\tilde{\psi}_i \, Q_i^{1+\zeta}}{1 - \iota_f(1-\delta_f)} \Big[1 + \frac{\delta_f \, \iota_f}{(1+\zeta)(1-\iota_f)} \Big] \tag{72}$$

The last form is in terms of observed rents and the calibrated $\tilde{\psi}_i$ of 70. At the baseline parameters the unimproved component is about 15 per cent of the total.

We thus see that the short run supply elasticity of construction is also the long run steady state supply elasticity of the stock of floorspace, and both are larger than the short run stock elasticity with respect to a price shock.

The steady state of the dynamic model is therefore exactly the static two-type model when the Gumbel scale is set so that $\varepsilon^g_{\mathrm{SS}}$ equals the static shape parameter ε,

$$\sigma^g = \frac{\iota}{\varepsilon\big[1 - \iota(1-\theta^g)\big]} \tag{73}$$

Under this nesting condition the static calibration recovers the steady-state fundamentals of the dynamic model directly, and the dynamic model adds only the parameters $\theta^g, \iota, \iota_f, \delta_f$; σ^g is implied rather than estimated. Away from the steady state the commuting matrix is not of the gravity form 68: it is a mixture of the stationary shares inherited by stayers and the choosers' shares 42, which respond to the whole future path.

C.8 Equilibrium solver

Given the initial conditions, the path of exogenous variables (total population and the exogenous public sector variables), and an assumption that we reach

steady state by period T, we can solve for a dynamic equilibrium path by finding the fixed point in rents and flow utilities that solves the dynamic equilibrium equations. First, we solve for the steady state value function variables: the residence values $\mathcal{R}^g_{iT}$, workplace values $\mathcal{F}^g_{jT}$, option values $\bar{\Omega}^g_T$ and floorspace values P_{iT} for steady state at the fixed level of the terminal exogenous variables. Then, starting with a guess of the path of Q_{it}, $x^g_{it} = ln(B^g_{it}) - (1-\beta^g)Q_{it}$, $y^g_{jt} = \ln w^g_{jt}$, we solve the value recurtions backward from T for the path of the residence values $\mathcal{R}^g_{it}$, workplace values $\mathcal{F}^g_{jt}$ and floorspace values P_{it}. Solving forward from the initial conditions, this sequence of value function variables determines the migration and construction decisions and therefore the path of the state variables: L^g_{ijt}, H_{it}. Given the state variables, we can solve for the within period prices and flow utility Q_{it}, $x_{it} = ln(B^g_{it}) - (1-\beta^g)$, $y_{jt} = w_{jt}$. By comparing these with the initial guess for the within period variables, we get a new guess to take forward for another pass of the loop, and repeat until we reach convergence. For our initial guess of the within period variables, we solve forward from the initial conditions using the same model but with myopic expecations (ie in making their migration and construction decisions workers and landlords incorrectly assume today's rents, wages and amenities will persist forever). The full dynamic equilibrium path can be solved in a few minutes on a laptop. The memory requirement is feasible because the separability of the value function means we never have to store a full NxNxT path of residence by workplace value functions. Similarly, the calculation is fast because there is no NxN variable that changes over time. For each loop we must calculate:

- Backward pass for value functions: the $2 \times N \times T$ set of residence and workplace value functions $\mathcal{R}^g_{it}$ and $\mathcal{F}^g_{jt}$ for each group, the single denominator of the movers choice shares Φ^g_t for each group (a log sum of those over residence workplace pairs weighted by the time invariant transport matrix) for each group, and the $N \times T$ floorspace values P_{it}.

- Forward pass for state variables: the $N \times T$ residence and $N \times T$ workplace populations L^g_{Rit}, L^g_{Fjt} for each group and the $N \times T$ floorspace stocks H_{it}. While the complete population state is $N \times N$, we never need to store it as we can calculate it in place using the resident and workplace vectors multiplied appropriately by the time invariant transport matrix.

- Within period: the NxT flow components x^g_{it}, y^g_{jt} for each group and the rental prices Q_{it}

All these variables have explicit closed form expressions, except the rental price which requires a within loop fixed point (and further speedup is gained by taking only one step of this fixed point each loop rather than going to convergence).

$$
\begin{aligned}
\mathcal{R}^g_{it} &= x^g_{it} + \iota(1-\theta^g)\,\mathcal{R}^g_{it+1},\\
\mathcal{F}^g_{jt} &= y^g_{jt} + \iota(1-\theta^g)\,\mathcal{F}^g_{jt+1},\\
P_{it} &= \iota_f\big[Q_{it+1} + (1-\delta_f)P_{it+1}\big],\\
\bar\Omega^g_t &= \theta^g\sigma^g \ln \Phi^g_{t+1} + \iota\,\bar\Omega^g_{t+1}, \qquad \Phi^g_t = \sum_{i\in\mathcal{S}^g_R} e^{(\iota/\sigma^g)\mathcal{R}^g_{it}} \sum_{j\in\mathcal{S}^g_F} e^{-\nu^g_{\mathrm{dyn}}\tau_{ij}}\, e^{(\iota/\sigma^g)\mathcal{F}^g_{jt}},
\end{aligned}
\tag{74}
$$

for $t = T-1,\dots,0$; these are (45), (54) and the definition of $\bar\Omega^g_t$ written as recursions. The first three are N-vectors per period (per type for the first two) and are what the forward pass consumes; the fourth is a scalar per period, cancels from every chooser share, and is computed only because welfare needs it. Φ^g_t is one dense matrix–vector product per type per period against the fixed commute cost matrix $e^{-\nu^g_{\mathrm{dyn}}\tau_{ij}}$, and it is the normalising constant of the chooser shares below, so it is formed once and shared.

Terminal values. At T the jump variables are set to their steady-state values at the terminal fundamentals,

$$
\begin{aligned}
\mathcal{R}^g_{iT} &= \frac{x^{g,\mathrm{ss}}_i}{1-\iota(1-\theta^g)}, \qquad & \mathcal{F}^g_{jT} &= \frac{y^{g,\mathrm{ss}}_j}{1-\iota(1-\theta^g)},\\
P_{iT} &= \frac{\iota_f\, Q^{\mathrm{ss}}_i}{1-\iota_f(1-\delta_f)}, \qquad & \bar\Omega^g_T &= \frac{\theta^g\sigma^g \ln \Phi^{g,\mathrm{ss}}}{1-\iota},
\end{aligned}
\tag{75}
$$

C.9 calibration

The dynamic model is calibrated at the 2016 Census, the last census before the relocation was announced, under the assumption that Sydney was then at, or close to, the steady state of the model with the 2016 fundamentals. This section first states what that assumption buys and what it costs, gives an outline of the whole procedure with the data each step needs, and then works through the steps in detail — for the single-shock model and for the separate-shock extension — indicating at each step what changes relative to the static two-type inversion, which is reproduced after the walk-through for reference.

What is held fixed. The location fundamentals a^g_j, b^g_i, ψ_i, the land areas K_i, the travel-time matrix τ_{ij}, the hazards and the preference and technology parameters are all time-invariant; the only things that move over the estimation window are the total populations, the public vectors and the policy instruments. Under this, the 2016 steady-state assumption does two jobs. It makes the observed commute-pair stock a gravity object — (68) in the single-shock model and its analogue in the extension — so that fundamentals can be inverted from marginals exactly as in the static model; without it the stock is a survival-weighted mixture of past choosers' flows and the residence term $\mathcal{R}^g_i$ cannot be read off the residence marginal. And it makes the observed floorspace stock

consistent with steady-state construction, $K_i J_i = \delta_f H_i$, which is what the land-value conversion and the recovery of the construction shifter rest on. Three relations are exact whether or not the city is in steady state — the choosers' gravity, within-period clearing $Q_{it} H_{it} = X_{it}$, and the developer's value recursion — and the checks at the end of this section use them to test the assumption rather than assume it.

Outline. In order:

1. *Parameters from the literature*, unchanged from the static model (Table 1): α, β^g, ζ and the spillover block. New to the dynamic model: the discount factors $\iota = \iota_f$ and the depreciation rate δ_f.

2. *Moving-opportunity rates.* Single shock: one hazard θ^g. Separate shocks: the residence and workplace hazards h_R^g, h_F^g and the joint probability θ_{RF}^g, from which the event probabilities follow. Data: the census five-year mobility tables for the residence hazard and the linked 2016–21 census place-of-work cross (ACLD) for the workplace hazard; the hazards are common across types in the baseline.

3. *Gravity*: the present-value commute decay ν_{dyn}^g from PPML on the 2016 SA2×SA2 commute matrix by education, with the movers' matrix as the assumption-free estimate and specification test.

4. *The Gumbel scale parameter*: calibrated in a similar manner to ε in the steady state Model; the other elasticities and κ^g follow from the hazards.

5. *Land values, rents and the construction technology*: Q_i from 2016 Valuer General land values, a steady-state relation; the flow shifter ψ_i from the steady-state shifter $\tilde{\psi}_i$; the initial stock from within-period clearing.

6. *The steady-state inversion*: wages from commute clearing, the cross-type scale, productivities, amenities, expenditure and $\tilde{\psi}_i$, then the agglomeration stripping — the static steps on the 2016 extracts, with two exponents changed under separate shocks.

C.10 Model against the evidence

Figures 40 and 41 compare the annual dynamic model with the synthetic-control estimates.

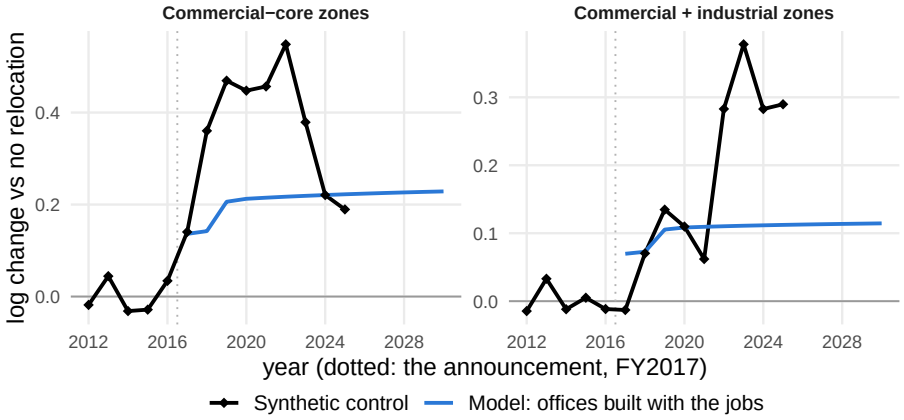


Figure 40: Land values: model against estimate. Log change in unimproved land value within 800 m of Parramatta Square, weighted by commercially zoned land, against the no-relocation path, on commercial-core zones (left) and commercial plus industrial zones (right). Black: the synthetic-control gap (the 15 most similar centres after the public-jobs and new-rail screens, the same parcels throughout, fitted in logs; onset 2017). Blue: the annual dynamic model, in which 13,884 public jobs move from the CBD to Parramatta Square over FY2018–22, each tranche anticipated from its announcement, with the offices built with the jobs. Dotted line: the announcement, FY2017.

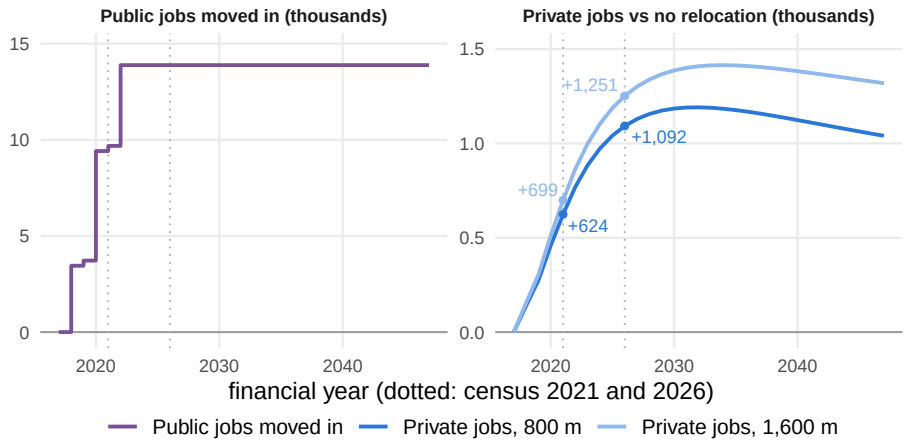


Figure 41: Employment in the annual dynamic model, offices built with the jobs. Left: public jobs moved in to Parramatta Square (thousands). Right: private jobs against the no-relocation path within 800 m and 1,600 m of Parramatta Square (thousands). Dotted lines: the 2021 and 2026 censuses. The decline after the early 2030s is the added office space depreciating without replacement while its occupants stay, a feature of how the experiment is set up rather than a result. The model's public sector is public administration (ANZSIC subdivision 75); the observed 2016–21 rise of 9,684 counts all government employment.