# References


Ahlfeldt, Gabriel M., and Elisabetta Pietrostefani. 2019. "The Economic Effects of Density: A Synthesis." *Journal of Urban Economics* 111 : 93–107. https://doi.org/10.1016/j.jue.2019.04.006.

Ahlfeldt, Gabriel M., Nathaniel Baum-Snow, and Remi Jedwab. 2023. *The Skyscraper Revolution: Global Economic Development and Land Savings.* Technical Report No. 18604.

Alonso, William. 1964. *Location and Land Use: Toward a General Theory of Land Rent.* Harvard University Press.

Balderston, K., K. Alferez, W. Uzzaman, and P. Garlick. 2025. *Housing Capacity Modelling Report: Housing Intensification and Resilience (Plan Change 120).* Technical Report.

Been, Vicki, Ingrid Gould Ellen, Michael Gedal, Edward Glaeser, and Brian J. McCabe. 2016. "Preserving History or Restricting Development? The Heterogeneous Effects of Historic Districts on Local Housing Markets in New York City." *Journal of Urban Economics* 92 : 16–30. https://doi.org/10.1016/j.jue.2015.12.002.

Bertaud, Alain, and Jan K. Brueckner. 2005. "Analyzing Building-Height Restrictions: Predicted Impacts and Welfare Costs." *Regional Science and Urban Economics* 35 (2): 109–25. https://doi.org/10.1016/j.regsciurbeco.2004.02.004.

Boeing, Geoff. 2017. "Osmnx: A Python Package to Work with Graph-Theoretic Openstreetmap Street Networks." *Journal of Open Source Software* 2 (12): 215. https://doi.org/10.21105/joss.00215.

Brueckner, Jan K. 2011. *Lectures on Urban Economics.* MIT Press.

Combes, Pierre-Philippe, and Laurent Gobillon. 2015. "The Empirics of Agglomeration Economies." In *Handbook of Regional and Urban Economics*, edited by Gilles Duranton, J. Vernon Henderson, and William C. Strange, vol. 5 of *Handbook of Regional and Urban Economics.* Elsevier. https://doi.org/10.1016/B978-0-444-59517-1.00005-2.

Cooper, G. C. K., and K. Namit. 2021. "City with a Billion Dollar View." *New Zealand Economic Papers* 55 (1). https://doi.org/10.1080/00779954.2018.1509110.

Cooper, Geoff, Ryan Greenaway-McGrevy, and James Allan Jones. 2026. "Measuring the Cost of Land Inputs to Housing Construction." *New Zealand Economic Papers* 60 (1): 1–23. https://doi.org/10.1080/00779954.2025.2536257.

Coulson, N. Edward, and Robin M. Leichenko. 2001. "The Internal and External Impact of Heritage Designations on Property Values." *Journal of Real Estate Finance and Economics* 23 (1): 113–24. https://doi.org/10.1023/A:1011120908836.

Donovan, Stuart, Thomas de Graaff, Arthur Grimes, Henri L. F. de Groot, and David C. Maré. 2022. "Cities with Forking Paths? Agglomeration Economies in New Zealand 1976–2018." *Regional Science and Urban Economics* 95 : 103799. https://doi.org/10.1016/j.regsciurbeco.2022.103799.

Donovan, Stuart, Thomas de Graaff, Henri L. F. de Groot, and Carl C. Koopmans. 2024. "Unraveling Urban Advantages—A Meta-Analysis of Agglomeration Economies." *Journal of Economic Surveys* 38 (1): 168–200. https://doi.org/10.1111/joes.12543.

Glaeser, Edward. 2008. *Cities, Agglomeration, And Spatial Equilibrium.* The Lindahl Lectures. Oxford University Press.

Greenaway-McGrevy, Ryan. 2025b. *Can Zoning Reform Increase Housing Construction? Evidence from Auckland.* Working Paper No. 17.

Greenaway-McGrevy, Ryan. 2025a. "Evaluating the Long-Run Effects of Zoning Reform on Urban Development, Housing Supply and Dwelling Prices." *Regional Science and Urban Economics* 111 : 103997. https://doi.org/10.1016/j.regsciurbeco.2024.103997.

Greenaway-McGrevy, Ryan, and James Allan Jones. 2024. "Mode Choice and the Effects of Rapid Transit Improvements on Private Vehicle Use and Urban Development." *Journal of Transport Economics and Policy* 58 : 1–36.

Greenaway-McGrevy, Ryan, and James Allan Jones. 2025. "The Welfare Effects of Character Protections on Neighbourhoods." *Economic Record* 101 (333): 203–32. https://doi.org/https://doi.org/10.1111/1475-4932.12851.

Hansen, Nikolaus. 2006. "The CMA Evolution Strategy: A Comparing Review." In *Studies in Fuzziness and Soft Computing*, vol. 192 of *Studies in Fuzziness and Soft Computing*.

Hansen, Nikolaus. 2016. "The CMA Evolution Strategy: A Tutorial.".

Hansen, Nikolaus, Youhei Akimoto, and Petr Baudis. 2019. "CMA-ES/pycma on Github." February. https://doi.org/10.5281/zenodo.2559634.

Hansen, Nikolaus, Raymond Ros, Nikolas Mauny, Marc Schoenauer, and Anne Auger. 2011. "Impacts of Invariance in Search: When CMA-ES and PSO Face Ill-Conditioned and Non-Separable Problems." *Applied Soft Computing Journal* 11 : 5755–69.

Hazledine, Tim, Stuart Donovan, and Christine Mak. 2017. "Urban Agglomeration Benefits from Public Transit Improvements: Extending and Implementing the Venables Model." *Research in Transportation Economics* 66 : 36–45. https://doi.org/https://doi.org/10.1016/j.retrec.2017.09.002.

Heintzelman, M. D., and J. A. Altieri. 2013. "Historic Preservation: Preserving Value?." *Journal of Real Estate Finance and Economics* 46 : 543–63. https://doi.org/10.1007/s11146-011-9338-8.

Jones, James A. 2021. "Essays in Quantitative Economics." Doctoral dissertation.

Kulish, Mariano, Anthony Richards, and Christian Gillitzer. 2012. "Urban Structure and Housing Prices: Some Evidence from Australian Cities." *Economic Record* 88 (282): 303–22. https://doi.org/10.1111/j.1475-4932.2012.00829.x.

Lees, Kirdan. 2014. "Quantifying the Costs of Land Use Regulation: Evidence from New Zealand." *New Zealand Economic Papers* 48 (1): 85–100. https://doi.org/10.1080/00779954.2013.854400.

Li, Zhenhua, Xi Lin, Qingfu Zhang, and Hailin Liu. 2020. "Evolution Strategies for Continuous Optimization: A Survey of the State-of-the-Art." *Swarm and Evolutionary Computation* 56.

Liotta, Céline, and co-authors. 2021. "Testing the Monocentric Standard Urban Model in a Global Sample of Cities." *Environment and Planning B: Urban Analytics and City Science*,. https://arxiv.org/pdf/2111.02112.pdf.

Mills, Edwin S. 1972. *Studies in the Structure of the Urban Economy.* Johns Hopkins University Press.

Muth, Richard F. 1969. *Cities and Housing: The Spatial Pattern of Urban Residential Land Use.* University of Chicago Press.

The New Zealand Initiative. 2025. *Competitive Urban Land Markets and the Planning Bill 2025.* Technical report. https://www.nzinitiative.org.nz/reports-and-media/reports/competitive-urban-land-markets-and-the-planning-bill-2025/document/915.

Noonan, Douglas S. 2007. "Finding an Impact of Preservation Policies: Price Effects of Historic Landmarks on Attached Homes in Chicago, 1990-1999." *Economic Development Quarterly* 21 (1): 17–33. https://doi.org/10.1177/0891242406296326.

Noonan, Douglas S., and Douglas J. Krupka. 2011. "Making-or Picking-Winners: Evidence of Internal and External Price Effects in Historic Preservation Policies." *Real Estate Economics* 39 (2): 379–407. https://doi.org/10.1111/j.1540-6229.2010.00293.x.

Nunns, Peter. 2015. *Topic 020 Viewshafts Expert Conferencing: Economic Costs of Viewshafts and Height Sensitive Areas.* Technical report.

NZ Transport Agency. 2013. *Economic Evaluation Manual.* Technical report. Wellington.

Stats NZ. 2025b. "Household Income and Housing-Cost Statistics: Year Ended June 2024." February. https://www.stats.govt.nz/information-releases/household-income-and-housing-cost-statistics-year-ended-june-2024/.

Stats NZ. 2025a. "Housing in Aotearoa New Zealand: 2025.". https://www.stats.govt.nz/.

Voronoi, Georgy. 1908. “Recherches Sur Les Paralléloèdres Primitifs. II. Nouvelles Applications Des Paramètres Continus à La Théorie Des Formes Quadratiques.” *Journal Für Die Reine Und Angewandte Mathematik* 134 : 198–287.

# 9. Technical appendix

This appendix provides the mathematical specification of the Alonso-Muth-Mills spatial equilibrium model used in this analysis, as well as details of the calibration algorithm and the computational implementation.

## 9.1. Mathematical framework

### 9.1.1. Utility and housing demand

Households have Cobb-Douglas preferences over housing ($q$) and a composite consumption good ($c$):

$$U = q^{\alpha} c^{1-\alpha} \tag{3}$$

where $\alpha$ is the share of expenditure on housing. A household at distance $x$ from the CBD earns a wage $W$ and incurs a commuting cost $tx$, giving a budget constraint:

$$p(x)q + c = W - tx \tag{4}$$

where $p(x)$ is the price of housing per unit of floor area at distance $x$. In equilibrium, all households achieve the same utility level $\overline{U}$, which implies a house price function that declines with distance:

$$p(x) = \left(\alpha^{\alpha}(1-\alpha)^{1-\alpha} \cdot \frac{W - tx}{\overline{U}}\right)^{1/\alpha} \tag{5}$$

### 9.1.2. Housing production

Housing is produced using land and capital in a Cobb-Douglas production function. The floor area ratio (FAR), denoted $h$, represents the intensity of development—the ratio of total floor area to land area. In the absence of constraints, profit maximisation yields:

$$h(x) = \left(\gamma \frac{p(x)}{p_K}\right)^{\gamma/(1-\gamma)} \cdot A^{1/(1-\gamma)} \tag{6}$$

where $\gamma$ is the capital share in production, $p_K$ is the cost of capital, and $A$ is total factor productivity in housing construction. Higher house prices thus lead to taller buildings, as developers substitute capital for land in the production of floor area.

### 9.1.3. Zoning and viewshaft constraints

In practice, development intensity is constrained by zoning regulations and, in some locations, by viewshaft height limits. Let $h_z$ denote the FAR cap for zone $z$. A viewshaft may impose a further constraint such that the effective FAR cap for a parcel in zone $z$ under viewshaft $v$ is:

$$h_{\text{eff}} = \min(h_z, h_v(x, y)) \tag{7}$$

where $h_v(x, y)$ is the FAR implied by the viewshaft height limit at that location. When the viewshaft binds ($h_v < h_z$), the development response is capped at the viewshaft limit. When it does not bind, the zone limit prevails.

For each zone $z$, there exists a **binding radius** $\hat{x}_z$ such that for distances $x < \hat{x}_z$ the FAR constraint is binding (the market-optimal FAR exceeds the cap), and for $x > \hat{x}_z$ the constraint is slack (the cap exceeds what the market would build). The binding radius is one of the key variables solved for during calibration.

### 9.1.4. Population distribution

The total population (measured in dwellings) is obtained by integrating housing consumption across the city. For each zone $z$, the population in the FAR-constrained region ($0 \leq x \leq \hat{x}_z$) and the free-market region ($\hat{x}_z < x \leq \overline{x}$) is:

$$N_{z,\text{FAR}} = \int_0^{\hat{x}_z} h_z \cdot J_1\left(\overline{U}\right) \cdot \text{area}(x)\, dx \tag{8}$$

$$N_{z,\text{free}} = \int_{\hat{x}_z}^{\overline{x}} h(x) \cdot J_2\left(\overline{U}\right) \cdot \text{area}(x)\, dx \tag{9}$$

where $J_1$ and $J_2$ are integral constants that convert floor area to dwelling counts:

$$J_1\left(\overline{U}\right) = (1-\alpha)^{(1-\alpha)/\alpha} \cdot \overline{U}^{-1/\alpha} \tag{10}$$

$$J_2\left(\overline{U}\right) = (1-\alpha)^{(1-\alpha)/(\alpha(1-\gamma))} \cdot \overline{U}^{-1/(\alpha(1-\gamma))} \cdot (\alpha\gamma/p_K)^{\gamma/(1-\gamma)} \cdot A^{1/(1-\gamma)} \tag{11}$$

The equilibrium requires that total population across all zones and distances equals the target population $N$:

$$\sum_z \left(N_{z,\text{FAR}} + N_{z,\text{free}}\right) = N \tag{12}$$

### 9.1.5. City boundary

The city extends to a boundary radius $\overline{x}$ where urban land rent equals the rural (agricultural) land rent $r_a$. The equilibrium utility level encodes this boundary condition:

$$\overline{U}(\overline{x}) = \frac{\alpha^\alpha \cdot \gamma^{\alpha\gamma} \cdot A^\alpha \cdot (1-\gamma)^{\alpha(1-\gamma)} \cdot (1-\alpha)^{1-\alpha} \cdot (W - t\overline{x})}{r_a^{\alpha(1-\gamma)} \cdot p_K^{\alpha\gamma}} \tag{13}$$

Higher rural land rents push the boundary inward (a more compact city with greater density), while lower rural rents allow sprawl. The city boundary $\overline{x}$ is one of the primary variables solved for during calibration.

### 9.1.6. Land rent

Land rent at distance $x$ is derived from the housing market equilibrium as the residual after capital costs:

$$r(x) = p(x) \cdot h(x) - p_K \cdot k(x) \tag{14}$$

where $k(x)$ is the capital input per unit of land. At the city boundary ($x = \overline{x}$), the equilibrium condition requires $r(\overline{x}) = r_a$. Land rent is converted to a capital value per square metre using a discount rate for comparison with observed land values from the DVR data.

## 9.2. Calibration

### 9.2.1. Parameter selection

Calibrating a spatial equilibrium model requires selecting parameters that govern household behaviour, production efficiency, and urban geography. In each case, the literature provides a range of estimates; our approach is to set bounds drawn from the literature and allow the optimisation to search within those bounds for values of $\alpha$, $\gamma$, $t$, $A$, and $r_a$ that, together with the spatial structure, reproduce Auckland's observed land values, city size, and built intensity. The resulting calibrated parameter set represents one internally consistent configuration that fits Auckland's 2024 reality.

**Housing expenditure share ($\alpha$).** The parameter $\alpha$ represents the proportion of household expenditure devoted to housing services. In the New Zealand context, Lees (2014) uses a value of 0.175, while other studies suggest a broader range of 0.25–0.325 to account for the high cost of housing in international metropolitan areas Greenaway-McGrevy and Jones (2024). However, December 2025 data from REINZ and StatsNZ shows Auckland renters spending an average of 40% of income on rent. Given that the model's wage gradient is driven in theory by the rental price of housing (where owners are both consumers and producers of housing), we set our optimiser upper bound for $\alpha$ = 0.40, and calibration selected this upper bound as the best fit to observed land values and house prices.

**Capital share of construction ($\gamma$).** The parameter $\gamma$ measures the elasticity of housing floor space with respect to capital inputs, reflecting building height and density. Standard benchmarks in the literature, including Bertaud and Brueckner (2005) and Kulish et al. (2012), frequently set $\gamma$ at 0.6. However, regional variations are significant; research on housing construction in France has estimated this share as high as 0.8, while other models explore sensitivities as low as 0.4 (Greenaway-McGrevy and Jones 2024). Higher values of $\gamma$ imply that land is more productive, meaning that development restrictions in such environments impose larger welfare losses. Our bounds for $\gamma$ were set at 0.4 and 0.8, and the optimiser selected a value of 0.539.

**Rural land rent ($r_a$).** The rural land rent defines the city edge, where urban land value equals its best alternative use. In Auckland, researchers have estimated this by taking average farm sales prices in surrounding regions such as Northland and Waikato, yielding a market value of roughly $20,000 per hectare. Applying a 4% capitalisation rate, this equates to an annual rental price of approximately $80,000 per km². However, farmland far from Auckland's urban areas is a poor proxy for the opportunity cost of urban land at the city's edge; lifestyle blocks within the Auckland region, where prices incorporate the option value of potential rezoning to urban uses, provide a more appropriate benchmark. We used benchmark property sale prices for lifestyle blocks in Rodney, Franklin, Waitākere near the urban fringe, which showed a range equivalent to around $1 - $2.5 million per km² (the model's unit for land area). Calibration selected a value of about $1.33M.

**CBD wages ($W$) and transport costs ($t$).** Estimation of wages and transport costs relies on the observational equivalence between commuting costs and the wage gradient. We gave the optimiser a range of $127,500 to $157,171 based on the 2024 median and mean household incomes for the Auckland region (Stats NZ 2025b).[7] The calibrated wage is $154,000. The marginal cost of transport has been derived from regressing median wages per statistical area against drive times; Greenaway-McGrevy and Jones (2024) found a cost of approximately $404 in annual wages per minute of freeflow drive time per worker. Using an average peak-hour speed of 25 km/h, this translates to a commuting cost in the order of $970 per kilometre per worker per year. Kulish et al. (2012) use a time-cost based approach of 60% of the wage rate per hour of commuting, which for a 25 km/h speed and a wage of $154,000 would yield about $1,665 of time cost per household per km of distance from the centre. Added to about $613 per household in annual pecuniary costs at $0.5 per km, the Kulish et al approach would give a $t$ value of $2,278. To allow some room for error and for the parameterised model to reach a closer fit to observed land value patterns, we set our bounds for $t$ at $500 and $2,500 per km, and calibration selected a value of $2,500.

### 9.2.2. Open-city parameters: agglomeration and congestion

The closed-city model treats household count $N$ and the wage $W$ as fixed. The open-city extension relaxes this: population migrates in or out until household utility equals an exogenous reservation level, and wages respond endogenously to population via agglomeration externalities. This introduces two additional mechanisms—an agglomeration equation linking wages to city size, and a congestion function linking transport costs to commuter volume—each governed by parameters that must be selected with care.

**Agglomeration elasticity ($\delta$).** Wages are determined by $W(N) = BN^{\delta}$, where $B$ is a scale constant calibrated so that $W(N_0) = W_0$ at the baseline population, and $\delta$ governs how strongly productivity responds to city size. NZ Transport Agency (2013) uses a weighted average elasticity of 0.065 across all industries, with values around 0.087 for knowledge-intensive sectors. Internationally, the broader meta-analytic literature places urban agglomeration elasticities in the range of 0.02–0.10, with most estimates clustering between 0.04 and 0.08 (Ahlfeldt and Pietrostefani 2019; Combes and Gobillon 2015). Donovan et al. (2024) synthesise 6,684 elasticities from 294 global studies, finding a mean agglomeration elasticity of 0.046 that falls to a bias-corrected estimate of 0.031 after accounting for publication selection. Their 95% credibility interval of 0.015–0.039 represents the most rigorous international benchmark available.

We select $\delta$ = 0.035. This sits within the bias-corrected 95% credibility interval of Donovan et al. (2024) and aligns closely with their bias-corrected global mean of 0.031. We consider this appropriately conservative given the simplicity of our approach to agglomeration, which does not capture the spatial variation in benefits that would be relevant. It is lower than the New Zealand micro-data production elasticity of 0.053 from Donovan et al. (2022). It is also conservative in that we expect some degree of sorting to work in Auckland's favour under viewshaft liberalisation. Donovan et al. (2022) show that aggregate-data estimates (0.071) are roughly double the micro-data estimates (0.037), with the gap attributable to self-selection of more productive workers into larger, denser cities. Auckland, as New Zealand's only metropolitan area of global scale, is the primary beneficiary of this sorting channel. When housing supply constraints are relaxed, the city is likely to attract disproportionately skilled workers.

**Congestion function.** The open-city model makes transport costs endogenous via a Bureau of Public Roads (BPR) congestion function:

$$t(N) = t_f \left[ 1 + \gamma_c \left( \frac{N}{N_{\text{cap}}} \right)^{\beta_c} \right] \tag{15}$$

where $t_f$ is the free-flow travel cost per kilometre, $N_{\text{cap}}$ is the effective road network capacity, $\gamma_c$ is the congestion multiplier, and $\beta_c$ is the congestion exponent. Greenaway-McGrevy and Jones (2024) use the same BPR structure, assigning $\beta_c = 4$ for link-level travel times following transport engineering conventions. Our model applies the

[7] Note that in the model all households earn the same wage, but the wage gradient is defined in combination with transport costs

congestion function at a more aggregate, city-wide level where individual link bottlenecks are smoothed out, so we adopt $\beta_c = 2.0$—preserving the essential convexity while reflecting the averaging inherent in a metropolitan-scale model.

The congestion multiplier $\gamma_c = 0.15$ and the free-flow ratio are calibrated jointly so that at the baseline population $N_0$, the congested travel cost equals the target value, with free-flow travel cost set at $\frac{t}{1.2}$, implying that current congestion adds roughly 20% to free-flow costs. This is consistent with Auckland's conditions, where peak-hour speeds in the inner city of approximately 25 km/h represent a meaningful but not extreme degradation from free-flow, while off-peak speeds are much closer to free-flow. The network capacity $N_{\text{cap}}$ is derived algebraically from these targets to close the system.

**Interaction between agglomeration and congestion.** The two mechanisms create opposing feedback loops. When viewshaft removal increases housing capacity, the city becomes more attractive, drawing in-migration. Agglomeration amplifies this: more workers raise productivity, attracting further migrants. Congestion acts as a brake: more commuters raise transport costs, reducing attractiveness. In equilibrium, the city expands until the marginal migrant is indifferent between Auckland and outside options, with the new population determined by the balance of these two forces. Our selection of moderate values for both $\delta$ and $\beta_c$ ensures that neither channel dominates unrealistically and that the model produces population responses plausible for a city of Auckland's size and connectivity.

### 9.2.3. Solving algorithm

The model is calibrated using the CMA-ES (Covariance Matrix Adaptation Evolution Strategy) optimiser, a derivative-free global optimisation algorithm well suited to the non-convex, multi-modal objective landscape of this problem. CMA-ES is a population-based metaheuristic that adapts the covariance structure of its search distribution over successive generations, enabling efficient exploration of rugged fitness landscapes without requiring gradient information. See Jones (2021) for a detailed overview. Our implementation uses the `pycma` python package, which is developed and maintained by the algorithm's authors and provides a robust and flexible interface for configuring the optimisation process (Hansen et al. 2019).

### 9.2.4. Optimisation variables

The parameter vector $\eta$ consists of:

$$\eta = \left[\overline{x}, \hat{x}_1, \hat{x}_2, ..., \hat{x}_{n_Z}, A^*, \alpha^*, \gamma^*, r_a^*\right] \tag{16}$$

where:

- $\overline{x}$ is the city boundary radius
- $\hat{x}_z$ is the binding radius for each zone $z$ (the distance at which the zone's FAR cap ceases to bind)
- The starred economic parameters ($A$, $\alpha$, $\gamma$, $r_a$) are optionally included in the optimisation when calibration bounds are specified for them; otherwise they are held fixed at user-supplied values.

This flexibility allows the calibration to jointly fit the spatial structure of the city and the economic parameters governing production technology, preferences, and the boundary condition, subject to bounds that reflect empirically plausible ranges.

### 9.2.5. Residual function

The composite residual is a weighted sum of four components:

$$\mathcal{R} = w_1 R_1^2 + \sum_z R_{2,z}^2 + w_3 \sum_i R_{3,i}^2 + w_4 R_4^2 \tag{17}$$

where:

- $\boldsymbol{R_1}$ **(Population fit)**: The normalised deviation between estimated and target population:

$$R_1 = \frac{N_{\text{est}} - N}{\tau_N} \tag{18}$$

where $\tau_N$ is either the target $N$ (relative tolerance) or a fixed tolerance in dwellings (absolute, e.g. 20,000). Absolute tolerance provides more intuitive control over how tightly the population target is enforced relative to other objectives when $N$ is large.

- $R_{2,z}$ **(FAR fit per zone)**: A symmetric residual for each zone comparing the model's estimated FAR at the binding radius to the zone's FAR cap:

$$R_{2,z} = \frac{h_{\text{est}}[z] - h_z}{h_{\text{est}}}[z] + \frac{h_{\text{est}}[z] - h_z}{h_z} \quad (19)$$

This penalises both overshooting and undershooting the zone FAR caps.

- $R_{3,i}$ **(Land value gradient)**: The relative deviation between model-predicted and observed land values at a set of distance nodes:

$$R_{3,i} = \frac{r_{\text{model}}(x_i) - r_{\text{obs}}(x_i)}{r_{\text{obs}}}(x_i) \quad (20)$$

The observed land prices are computed as distance-binned medians of land value per square metre from the DVR data.

- $R_4$ **(House price level)**: The relative deviation between the model-predicted house price and a target price at a reference distance:

$$R_4 = \frac{p_{\text{model}}(x^*) - p_{\text{target}}}{p_{\text{target}}} \quad (21)$$

Each component can be individually weighted (or disabled by setting its weight to `None`) to control the balance between fitting the population distribution, the built intensity pattern, the land value gradient, and the house price level.

### 9.2.6. Parameter interactions

Several key interactions between parameters are worth noting, as they inform the calibration strategy:

- **$A$ and $\gamma$ trade off**: Increasing either parameter raises the housing supply capacity of the model (helping fit the population target) but also tends to lower house prices.
- **$r_a$ resolves the population-price tension**: Higher rural land rents simultaneously raise land prices and increase population density (by compressing the city), making $r_a$ the key lever for jointly fitting population and price targets.[8]
- **$\alpha$ controls dwelling size**: Lower $\alpha$ means smaller dwellings (less housing per unit of income devoted to housing) and a steeper house price gradient with distance.
- **$\gamma$ controls supply elasticity**: Higher $\gamma$ means more responsive FAR to price changes, allowing the model to replicate the wide range of built intensities observed across Auckland's zones.

## 9.3. Model implementation

The model is implemented in Python (version 3.11+) in the `viewshafts` package, a proprietary library developed by Crow Advisory. The core spatial equilibrium calculations are fully vectorised using NumPy for performance.

The key software components are:

- **`AnnularAMM`**: The main orchestrator class. Manages data preparation (extracting parcels, applying FAR limits, filtering by viewshaft scenario), parameter configuration, and CMA-ES optimisation. Exposes a `solve()` method that returns a `CalibrationResult` containing the optimised parameters, population matrix, and utility level.
- **`ViewshaftScenario`**: Manages scenario definition by filtering viewshafts by maunga name or schedule code. Handles HSA associations and applies an area threshold (default 50%) for determining whether a parcel is sufficiently within a viewshaft to be constrained. Returns the binding height limit per parcel for the scenario.
- **`FARScenario`**: Combines zone-based FAR limits with viewshaft-derived height limits into a single effective FAR constraint per parcel, using the more restrictive of the two where both apply.
- **`DataLoader`**: A lazy-loading cache for all geospatial datasets, including LINZ parcels, Stats NZ SA2 boundaries, viewshaft overlay polygons with plane coefficients, and the OpenStreetMap road network graph.

[8] The recommended empirical range for $r_a$ in Auckland is \$150,000 to \$2,500,000 per km² per year, based on pastoral, peri-urban, and lifestyle land sale values.

- **`as_calibrate_fn`**: The core calibration function. For a given parameter vector $\beta$, it computes the equilibrium utility, estimates FAR at each binding radius, integrates population across FAR-constrained and free-market annuli for all zones, and evaluates the composite residual (Equation 17). Supports both scalar residual output (for optimisation) and full population matrix output (for analysis).
- **`CalibrationResult`**: A dataclass holding the optimised $\beta$ vector, residual value, calibrated economic parameters, FAR vector, distance nodes, equilibrium utility, and the full population matrix. Includes plotting methods for diagnostic visualisation.

## 9.4. CMA-ES configuration

The CMA-ES optimiser is configured as follows:

**Table 9: CMA-ES optimiser configuration**

| Setting | Value |
|---|---|
| Population size (of candidate solutions) | 100–2,500 (auto-scaled by dimension) |
| Maximum iterations | 500–2,000 |
| Function tolerance | $10^{-11}$ |
| Spatial variable bounds | $\left[D_{\text{Secs}[1]}, \infty\right)$ |
| Economic parameter bounds | Per-parameter (e.g. $\alpha \in [0.2, 0.4]$) |
| Initial step sizes | Tighter for economic parameters based on bound width |

When economic parameters are included in the optimisation, each receives per-variable bounds and tighter initial step sizes (CMA standard deviations) relative to the spatial parameters. After optimisation, the economic parameter values are extracted from $\beta$ and stored as an updated `AmmParams` instance on the result.

## 9.5. Open-city extension

The open-city version of the model allows population $N$ to adjust endogenously. Wages are determined by an agglomeration function:

$$W(N) = BN^{\delta} \tag{22}$$

where $\delta$ = 0.035 is the agglomeration elasticity and $B$ is a scale constant computed so that $W(N_0) = W_0$ at the calibrated closed-city population $N_0$. This ensures that the closed- and open-city models agree at the baseline.

Commuting costs are also population-dependent through a congestion function. The effective per-kilometre transport cost is:

$$t(N) = t_0 \left(1 + \gamma_c \left(\frac{N}{N_{\text{cap}}}\right)^{\beta_c}\right) \tag{23}$$

where $t_0$ is the free-flow transport cost, $N_{\text{cap}}$ is a capacity parameter, and $\gamma_c$ = 0.15 and $\beta_c$ = 2 are the congestion elasticity and exponent respectively. This ensures that the benefits of agglomeration are partially offset by increased travel times as the city grows.

The open-city equilibrium pins utility to the closed-city baseline $\overline{U}_0$ (representing the national reservation level). The city boundary is obtained by inverting the utility expression (Equation 13) for $\overline{x}$:

$$\overline{x}(N) = \frac{W(N) - \overline{U}_0 \cdot r_a^{\alpha(1-\gamma)} \cdot p_K^{\alpha\gamma}}{t(N) \cdot \alpha^{\alpha} \cdot \gamma^{\alpha\gamma} \cdot A^{\alpha} \cdot (1-\gamma)^{\alpha(1-\gamma)} \cdot (1-\alpha)^{1-\alpha}} \tag{24}$$

For a given $N$, the model computes $W(N)$, $t(N)$, and $\overline{x}(N)$, then solves for the spatial equilibrium as in the closed-city case. The equilibrium population is found by searching over $N$ until the total population implied by the spatial equilibrium matches $N$—that is, the fixed-point condition:

$$\sum_z (N_{z,\mathrm{FAR}} + N_{z,\mathrm{free}}) = N \tag{25}$$

is satisfied simultaneously with the utility constraint. In practice, this is solved using a bounded scalar optimisation over $N$.

## 9.6. Data processing pipeline

The full pipeline from raw data to model inputs involves the following stages:

1. **Viewshaft parsing**: XML ⟶ vertex sets ⟶ polygon GeoDataFrames ⟶ plane equation coefficients ($a$, $b$, $c$, $d$) per viewshaft ⟶ parcel matching.
2. **Height limit assignment**: Contour sampling ⟶ spatial join ⟶ plane evaluation (all heights, all contour points) ⟶ ground elevation subtraction ⟶ parcel-level height limit stack.
3. **FAR computation**: Zone FAR caps ∩ viewshaft height-derived FARs ⟶ effective FAR per parcel.
4. **Distance computation**: OSM network ⟶ Dijkstra shortest paths ⟶ parcel-level distances to CBD ⟶ travel-time matrix ⟶ "minutes-to-centre" bins.
5. **Zone-distance structure**: Zone-distance bin aggregations ⟶ cumulative zone proportions ⟶ $\theta$ vector ⟶ $\boldsymbol{H}_{\mathrm{SecZ}}$ matrix.
6. **Calibration targets**: DVR land values ⟶ distance-binned medians ⟶ observed price gradient.