Can the Urban Poor Avoid Flood Risks? The Case of Cape Town, South Africa*

Paolo Avner, Charlotte Liotta, Thomas Monnier, Claus Rabe, Harris Selod†

September 2026

Abstract

In low- and middle-income country cities, poor households often reside in hazardous locations, including flood-prone areas. This can be due to poor information about flood risks or acceptance of these risks in the face of lower housing prices. Poor households are also more vulnerable to floods than richer households given the low-quality housing they occupy. Does information on flood risks help households make better location and housing choices? To what extent will these choices be revised with increased flood risks due to climate change? To answer these questions, we build a polycentric urban economics model with heterogeneous income groups, formal and informal housing, and flood risks. The model is calibrated to Cape Town (South Africa) and simulations are run to assess the impact of flood risks on land values and income segregation within the city, distinguishing between the effects of three types of floods (fluvial, pluvial, and coastal). Although total damages from floods are greater for rich households, they represent a larger relative share of poor households' incomes. Better information encourages the partial adaptation of poor households mainly through relocation, and especially in the face of fluvial risks, reducing up to 40% of expected damages. Added risks from climate change undo these gains by roughly 10 percentage points.

Keywords: floods, climate change, adaptation, resilience, informal housing
JEL Classification: Q51, Q54, R14, R31

*We thank the City of Cape Town for its support and facilitation of access to the data used in this paper. We are especially grateful to Amrish Bisoon, Hugh Cole, Paul Court, Robert McGaffin, Elmarie Nel, Jodie Posen, Kristoff Potgieter, Carol Wright, and all participants to the World Bank technical assistance organized as part of the Resilience Pillar of the Cities Support Program led by the National Treasury of the Government of South Africa. We are also indebted to Haroon Bhorat, Lukasz Grzybowski, Andrew Kerr, Martine Visser, and other people at the University of Cape Town for their valuable feedback. We would also like to thank Jun Goto, Seiro Ito, Kelsey Jack, Minoru Osawa, and Sen Xue for their comments. Funding from the Swiss State Secretariat for Economic Affairs is gratefully acknowledged.

†Avner: World Bank, GFDRR, pavner@worldbank.org; Liotta: Universitat Autònoma de Barcelona, ICTA, charlotte.liotta@uab.cat; Monnier: Hitotsubashi University, Institute for Advanced Study, thomas.monnier@r.hit-u.ac.jp; Rabe: Palmer Development Group, claus@pdg.co.za; Selod: World Bank, DEC, hselod@worldbank.org

1 Introduction

In low- and middle-income countries, urban growth can be accompanied by greater vulnerability of poor households who tend to settle in flood-prone areas. The World Bank estimates that 1.8 billion people currently live in flood-prone areas, including 600 million people in cities (World Bank, 2026). The phenomenon is on the rise in all regions of the world as informal settlement growth is occurring at a higher rate in flood-prone than in flood-safe areas (Rentschler et al., 2023). In South African cities, up to 19 percent of the population could be exposed to flood risks by 2050 (World Bank, 2022).

In this paper, we propose the first urban simulation model that simultaneously accounts for all three types of flood risks (fluvial, pluvial, and coastal). Accounting for the whole range of flood risks makes it possible to assess how they differentially impact cities, and is key to designing better targeted policy responses.

As our model ambitions to offer a realistic representation of a developing country city, it is important to account for key features such as income heterogeneity and the coexistence of formal and informal housing (including both squatter settlements and structures informally erected in the backyard of public housing units). Building on a previous version of the model developed for Cape Town, South Africa (Monnier et al., 2026), we overlay a discrete two-dimensional version of the standard urban monocentric model (Fujita, 1989) on a grid of pixels that capture the city's geographic features. We then simulate the construction decisions of developers along with the housing and location choices of households who compete spatially to access multiple employment centers. At the same time, we account for land-use and building regulations, transport networks, natural constraints, and exogenous amenities. A simplified version of this modeling approach has proven effective in predicting urban spatial structures across the world (Liotta et al., 2022).

We add flood risks to the framework by considering the probability of flood occurrence in each grid cell, and the corresponding damages caused to housing structures and contents, taking into account the lower quality of informal buildings. This focuses on the most immediate channels through which flood risks affect city patterns over the long term (Merz et al., 2010; Pharoah, 2014; Mosimann et al., 2018).[1] Agents may internalize these risks if they consider the expected annual damages (based on probabilistic flood maps) as it affects the depreciation of their housing capital and the quantity of goods they consume (Huizinga et al., 2017). This extends the framework in Avner et al. (2022) by incorporating the impact of flood damages in both housing supply and demand.

We further distinguish between fluvial, pluvial, and coastal floods. Typically, fluvial floods are water overflows from rivers, whereas pluvial floods designate surface water floods or flash floods, caused by extreme rainfall independently of an overflowing water body. Coastal floods encompass storm surges, periodic tides, and gradual sea-level rise. In our model, flood risks

[1] In this paper, we do not focus on damages to infrastructure (Hallegatte et al., 2019), indirect effects on health (Picarelli et al., 2017; Paterson et al., 2018) or public services (Hammond et al., 2015; Allaire, 2018).

do not only affect population and structures through their direct exposure to floods, but also through spatial equilibrium effects (i.e., given endogenous housing prices and spatial sorting of income groups across the city).

Our analysis is based on the comparison of various simulations of the model. To assess the role of information on flood damages, we compare equilibria under scenarios with and without flood risk anticipation. To assess the impact of climate change, we then compare equilibria with flood risk anticipation under a scenario where flood risks remain at their current level and another scenario where flood risks are amplified. We find that although total damages from floods are greater for rich households (who consume more housing), they represent a larger relative share of poor households' incomes (who therefore tend to react more strongly to flood risks). Better information on flood risks encourages their relocation outside flood-prone zones and lowers their willingness to pay for exposed housing. However, because affordable flood-safe areas may also be located far from job centers, the poor do not entirely move away from flood risks. This underlines a fundamental trade-off poor households face between accessing economic opportunities and accepting exposure to flood risks. When households are informed about flood risks and when these risks increase, most of the changes in equilibrium involve households avoiding localized fluvial flood zones that face severe destruction risks. On the contrary, pluvial flood zones are typically more widespread, with lower destruction risks. This also explains why, in the face of climate change, poor households are willing and able to reduce most of the adverse consequences of fluvial, but not of pluvial floods. Coastal flood risks mostly affect richer households who are less willing to adapt, all the more so as these risks are less severe than fluvial/pluvial risks.

This paper fits in a tradition of land use and transport integrated models (LUTI) applied to the evaluation of urban climate policies in various economic contexts. For instance, in the context of Paris, France, Viguié and Hallegatte (2012) look at trade-offs and interactions between urban policies, including a policy that prohibits new buildings in flood-prone areas (flood-risk zoning). In the context of Cape Town, Liotta et al. (2024) study the impact of a fuel tax on spatial inequalities. They find that the poorest households, living in informal settlements or in subsidized housing, have few or no ways to adapt to increases in fuel prices by changing housing type, adjusting their location or housing consumption, or shifting transportation modes. In the wake of the above literature, the current paper focuses on adaptation to climate change and studies more specifically the impact of flood awareness policies on spatial inequalities, a perspective that was previously overlooked.[2]

Our approach also relates to the"quantitative spatial model" (QSM) literature that emerged over the past decade (Redding and Rossi-Hansberg, 2017) and is increasingly applied to emerging country contexts (Sturm et al., 2023). Apart from the fact that we allow for an endogenous city fringe (which should be accounted for in contexts of sprawling and growing cities as in South Africa), the main conceptual difference with the QSM approach is that we consider job locations

[2]See Carleton et al. (2024) for a systematic review of climate change adaptation policies.

and wages as exogenous. We do not consider this simplification to be a strong limitation for our purpose given that floods generally do not lead to massive displacements of economic activity (Kocornik-Mina et al., 2020).[3] Besides, it allows for a more tractable model, with typically more dimensions of heterogeneity and a more granular geography. To the best of our knowledge, our model is also the first to allow for a realistic urban simulation that integrates all types of flood risks within an internal city structure. Although some quantitative spatial models address coastal risks (Lin et al., 2024; Balboni, 2025), these models focus on firms rather than on households. Furthermore, they consider sea-level rise only and do not account for other components of coastal risks such as tides and storm surges. The structural economic literature on fluvial and pluvial risks appears to be less extensive, likely due to the periodic rather than persistent nature of these phenomena, which makes them difficult to model and analyze. The rest of the literature on flood adaptation is mostly in reduced form and focuses exclusively on formal housing markets (Ortega and Taşpınar, 2018; Muller and Hopkins, 2019; Ellen and Meltzer, 2024). A notable exception is Varela Varela (2023), which uses a residential segregation model to explain heterogeneous price recoveries following a hurricane.

The outline of the paper is as follows: Section 2 presents stylized facts regarding flood risks in Cape Town. Section 3 describes the data. Section 4 details the model and Section 5 explains how it is estimated. Section 6 presents the results for our two sets of simulation comparisons: the first one introduces information on flood risks, and the second one the effects of climate change. Section 7 introduces the next steps and Section 8 concludes.

2 Context

Cape Town, a sprawling city of four million inhabitants is located on a broad, sandy plain connecting the mountainous Cape peninsula to the mainland. It is among the cities most affected by fluvial and pluvial flood risks in Sub-Saharan Africa, with 160,000 households directly exposed to flood risks and average annual damages estimated at almost \$16.4 million (Gonzales et al., 2023). Damages from coastal floods are less important in absolute terms (around \$370,000 annually), but could sharply increase and affect more people with climate change and sea-level rise (Hallegatte et al., 2013). For reference, we show a map of Cape Town's major flood plains in Figure 1.

A large fraction of the population resides in the Cape Flats, a predominantly low-lying area of about 25x25km where approximately 200,000 poor households (15-20% of the population) live in informal settlements. The hydrology of the Cape Flats renders poor households living in informal settlements vulnerable to both fluvial and pluvial flooding. Fluvial flooding is episodic, typically more intense, and localized along Cape Town's four perennial rivers and stormwater

[3] Choosing between the two modeling approaches involves accepting either exogenous location-specific productivity differences, which capture the structural residuals left unexplained by QSMs, or exogenous wages in each employment center, as in our approach.. For an example of a QSM approach to Cape Town land use, see Cole et al. (2025).

Figure 1: Cape Town's major flood plains

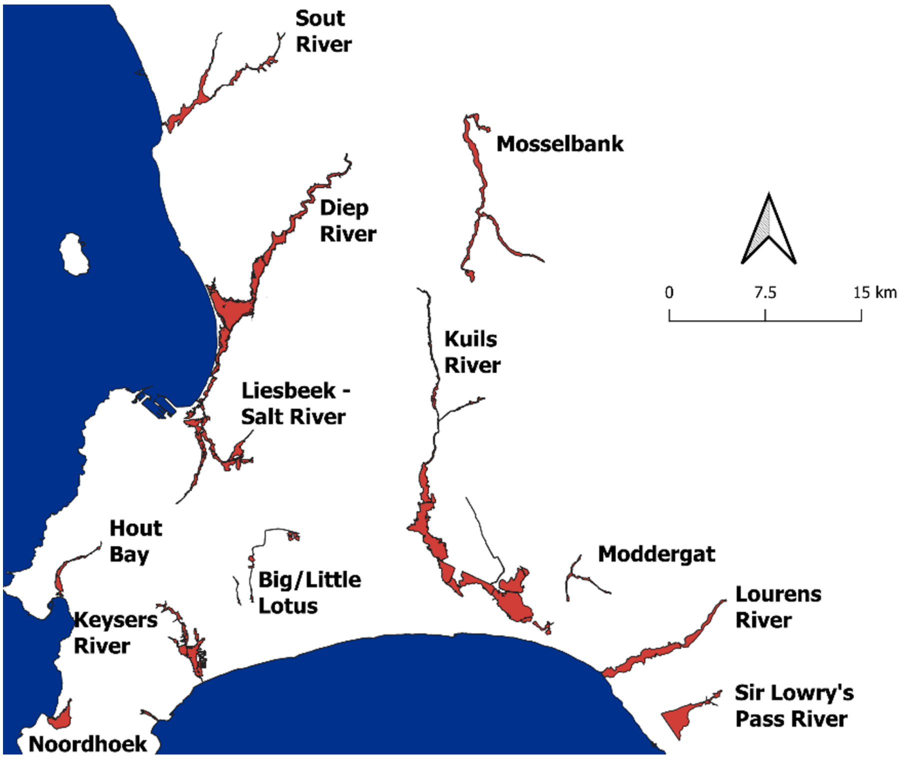

Source: City of Cape Town

channels. Pluvial flooding, on the other hand, is persistent and widespread across the Cape Flats area: the vast network of seasonal rivers, streams, and wetlands transecting the area is inundated for a few weeks annually during the rainy winter season. From there, drainage is encumbered by flat terrain and a water table lying between 1-3m from the surface in the dry summer months, but rises by between 1-2m during winter, causing "rising water" or "seepage" in topographic depressions. The impact of hydrology on households is aggravated by inadequate drainage infrastructure, which leads to localized ponding. Figure 2 shows a more detailed map of the Cape Flats and their exposure to floods in flood-prone depressions (green areas) and flood-prone riverines (blue areas). It also shows the location of informal settlements (dark grey and black areas) in proximity of these zones.

The magnitude of the flood impacts on poor households is determined both by nature and by human behavior. The main human cause of impact is the massive, informal occupation of large areas of flood-prone wetlands and detention ponds by poor households who build vulnerable structures from corrugated iron sheets. The propensity for households to locate in flood-prone areas is not incidental, but attributable to the fact that formal development is prohibited in flood-prone areas, leaving vacant remnants of land throughout the Cape Flats area. Furthermore, South African courts have pronounced that local authorities are not allowed to evict illegal land occupants from erected structures unless alternative accommodation is provided. Land occupation events typically occur in the run-up to winter since local authorities are reluctant to evict residents during winter on humanitarian grounds. In the remainder of this paper, we will

Figure 2: Flood-prone zones in the Cape Flats

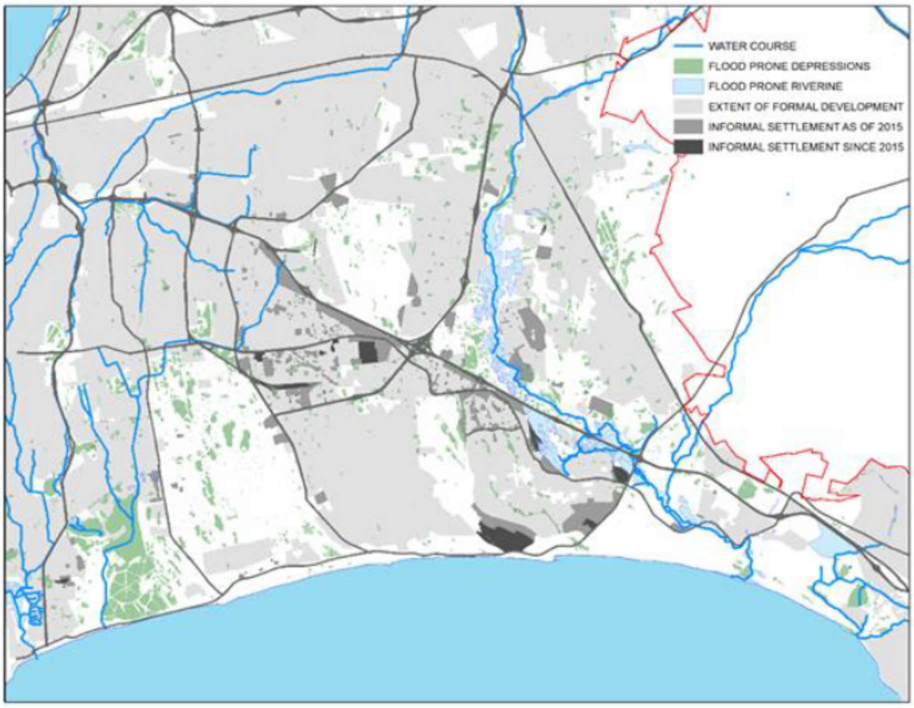

Source: City of Cape Town

consider that tolerated informal settlement zones are exogenously given.[4]

With population growth and given these ongoing spatial dynamics, the number of informal households in flood-prone areas has been multiplied by 6 from 1995 to 2020, while population density in these areas has increased by roughly 50% (based on aerial imagery). In what follows, we uncover the economic mechanisms behind such development, and assess how flood damages are likely to evolve following adaptation to better information and climate change.

3 Data

This section starts by presenting the data we use to estimate the model, and how we process it.

Grid We use a grid of 500x500m resolution that encompasses the whole urban area and corresponds to the grid that the City of Cape Town uses for planning purposes. All other datasets are either spatially aggregated or disaggregated to fit the grid.

Flood data We use the FATHOM (Sampson et al., 2015) and the Deltares (2021) data, two global spatial datasets of flood hazards. They provide flood water extent and depth for respectively pluvial/fluvial and coastal hazard scenarios, expressed as "return periods", which indicate the likely frequency of occurrence (i.e., once every number of years) for each flood event category. When overlapping, we consider the maximum flood depth across flood types as

[4]See Monnier et al. (2026) for a discussion of this assumption.

Figure 3: Expected capital depreciation rate due to floods (for housing contents only) and increase induced by climate change

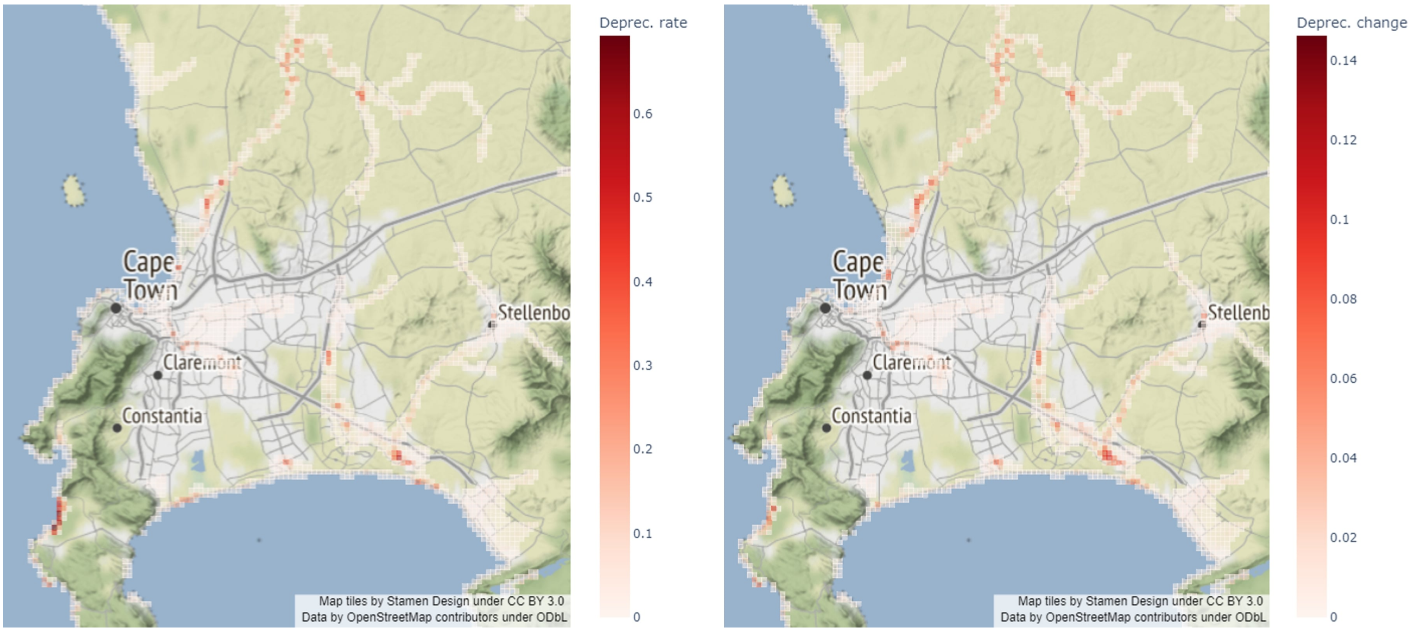

Note: The map on the left shows the values at the benchmark. The map on the right shows the incremental change (in absolute terms) when considering climate change.

water would typically flow to other areas instead of piling up. The data are at a 3 arc-second (approximately 90m) resolution.

Because our geographic units are 30 times larger than the units used in the raw flood maps, the flood data that we use is sufficiently granular to capture flood impacts at the grid cell level. Note however that we use the "undefended" versions of the datasets, i.e. the ones that do not account for infrastructure and other potential protection investments affecting flood hazards. This is because the global data is not precise enough to properly account for their spatial distribution at the local level. An exception we make is for the presence of drainage systems, which we assume to be linked with formal concrete housing structures: such dwelling units will not be affected by the least severe (but more frequent) flood events.

In every grid cell, the flood data is converted into an expected fraction of capital destroyed, collapsing all return periods into a single annualized value. To do so, we borrow "structural" damage functions by housing type from Englhardt et al. (2019) that convert flood depth into building destruction shares. We also consider damages in terms of "contents", expressed as the fraction of (semi-)durable goods that households consume and that is vulnerable to floods. To obtain the expected fraction of contents destroyed, we use the flood depth-damage function proposed by De Villiers et al. (2007).

Finally, we consider how climate change might affect flood hazards. For coastal flood risks, we rely on specific flood maps provided by Deltares, that are based upon the IPCC's RCP 8.5 scenario projected onto the year 2050. Since the FATHOM data does not come with similar maps for pluvial/fluvial flood risks, as a proof of concept, we simply multiply all annual probabilities by 2 (i.e., we divide return periods by 2). This increase, although arbitrary, is of the same order of magnitude as the one observed for coastal flood risks. It should be noted that those scenarios are generally thought of as being relatively pessimistic.

To illustrate how we account for flood risks, we represent in Figure 3 the spatial distribution of the content depreciation rate (which does not vary across housing types), and how it incrementally changes with increased flood risks due to climate change (see Section 4 for more details on structure and content depreciation rates). As expected, flood exposure follows the coastline (for coastal flood risks), waterways (for fluvial flood risks), and topological depression areas (for pluvial flood risks) in inhabited areas. The added risks from climate change are not negligible, reaching more than 10 percentage points in most exposed zones. Note that these maps only represent intermediate inputs, as equilibrium damages will depend on households' housing and location choices.

Socioeconomic data The spatial distribution of the population is taken from the 2011 National Census. As explained in Section 4 below, we consider four housing types, namely formal private housing (FP), formal subsidized housing (FS), informal settlements (IS), and informal backyards (IB). We also define four income groups of interest by choosing income-group thresholds such that only the lowest income group is eligible for subsidized housing programs, and so that the two highest income groups are not observed to reside informally (but are nevertheless distinguished to account for high income heterogeneity).[5]

We use the City of Cape Town's transport model to retrieve transport times and distances between residence and job locations for each transport mode, along with the estimated number of households per employment center and income group. The monetary transport costs are retrieved from additional external sources. We also use aggregate statistics on residence-workplace distances in Cape Town, derived from Cape Town's 2013 Transport Survey, to complete the estimation of households' expected income net of commuting costs (see Section 5).

Land availability is pre-determined for each housing type in each grid cell. Areas of subsidized housing are identified from the cadastre of the City of Cape Town. The area available for backyard housing is estimated as the yard size of these units (as in the context of Cape Town, most of so-called informal "backyarding" occurs in such precincts). Informal settlement areas are obtained from the characterization of enumeration areas (census blocks) in the 2011 Census: in the context of Cape Town, they correspond to peripheral, publicly-owned land originally reserved for future roads, social facilities, or public housing. It therefore makes sense to consider them as being exogenously given as we do. Their boundaries reflect the tolerance of authorities for these locations. Land available for formal private development corresponds to all land that is not constrained for housing construction (including exogenous commercial floor space). Additional height restrictions apply across the city.

[5]These groups reflect skill heterogeneity within the population. Workers employed in the same locations but belonging to different groups will therefore obtain a different wage.

The amenities that we consider include natural amenities (such as slope and proximity to the ocean) as well as urban amenities (such as the proximity to the historical center). We extract them from the City of Cape Town's open data portal, and consider them to be exogenous.[6] Finally, we use property prices extracted from the City of Cape Town's geocoded dataset on property transactions for 2011, as well as data on dwelling sizes made available to us by the City of Cape Town. These will serve to estimate the parameters of the housing production technology used by formal private developers.

4 Model

4.1 Main assumptions

We build on the model presented in Monnier et al. (2026) and augment it with flood risks to define our benchmark. As in their framework, we consider that each grid cell in the city is indexed with a vector of coordinates $x = (x_1, x_2)$ and has an exogenous quantity of available land for residential development per housing type $h = FP, FS, IS, IB$, $L_h(x)$, where FP refers to formal private housing, FS refers to formal subsidized housing, IS refers to informal settlements (i.e., traditional squatter settlements), and IB refers to informal backyard structures. We further split the population into four income groups (ranked from poorer to richer): only income group 1 has access to FS housing while income groups 3 and 4 are never observed to live in IS or IB housing.

In addition, each location is characterized by an exogenous amenity index $A(x)$ (centered around one), and an exogenous informal disamenity index B_h (comprised between zero and one) associated with informal housing types $h = IS, IB$. As both indices will enter the utility function of residents as multiplicative terms, the lower their value, the greater the disamenity associated with location or informal housing. Note that informal neighborhoods do not only differ from formal ones by (horizontally dense) residential land use and (poorer) construction quality, but also in terms of public service (lower) coverage and (higher) subsidies. This will implicity affect the housing demand elasticity across housing types through the term B_h.

In this framework, households compete with one another for locations and housing types, anticipating their probability to work in the different employment centers given transport costs, wages and each income group's probability of being employed. This is inclusive of informality in the labour market, and the probability to be unemployed is also accounted for. Equilibrium will be defined "ex ante", i.e. before worker's specific employment centers are known, but accounting for the probability to be working in each one of them.

Households Households choose their consumption of composite good z and housing quantity q, along with their residential location x and housing type h, by maximizing Stone-Geary

[6] In theory, this leads to a worse fit than recovering amenities as model residuals (as in QSMs). However, in our case, we find that this does not lead to strong underfitting.

preferences:

$$U(z,q,x,h) = z^{\alpha}(q-q_0)^{1-\alpha}A(x)B_h \quad (1)$$

where $0 < \alpha < 1$ is the composite good elasticity (or budget share), $1-\alpha$ is the surplus housing elasticity, and $q_0 > 0$ is the basic housing need.

In doing so, they face the following budget constraint, which is specific to each income group $i = 1,..,4$:

$$\begin{aligned}&\tilde{y}_i(x) + \mathbb{1}_{\{h=FS\}}\mu(x)YR_{IB}(x)\\ &\quad= \left(1+\gamma\rho^{content}(x)\right)z + q_hR_h(x) + \mathbb{1}_{\{h\neq FP\}}\left(\rho + \rho_h^{struct}(x) + \mathbb{1}_{\{h\neq FS\}}\delta\right)v_h \quad (2)\end{aligned}$$

In equation (2), the left-hand side measures revenues, and the right-hand side breaks down expenses. On the left-hand side, $\tilde{y}_i(x)$ is the expected income net of commuting costs for a household of income group i living in location x: since households choose their residential location before finding a job, they consider the likelihood that they will find employment (or not) in the different job centers, along with the associated commuting costs from each residential location and the (exogenous) income they would obtain given their income group at each specific workplace.[7]

When living in formal subsidized housing ($\mathbb{1}_{\{h=FS\}}$), households benefit from a backyard of fixed size Y and have the additional possibility to rent out an endogenous fraction $\mu(x)$ of that backyard at the endogenous market rent $R_{IB}(x)$ (as in Brueckner et al. (2019)), which adds up to their wage income.

On the right-hand side of equation (2), households need to pay for composite good z (whose price is normalized to one). γ is the fraction of (semi-)durable goods that is exposed to floods and that depreciates at rate $\rho^{content}(x)$. This corresponds to the expected fraction of capital destroyed in each location due to floods as shown in Figure 3.

Households also need to pay for housing quantity q_h at the endogenous annual rent $R_h(x)$. Regarding housing quantities, although housing consumption is a choice variable for $h = FP$, it is fixed for $h = FS, IS, IB$ given that formal subsidized housing and informal housing structures are both standardized in South Africa. Regarding "rents", note that our setting captures different situations across housing submarkets: without loss of generality, formal private residents are assimilated to tenants who rent their housing from developers who themselves previously bought land from absentee landlords (see the "housing supply" subsection below). Whereas "backyarders" pay rent to the beneficiaries of formal subsidized housing, informal settlement dwellers pay a fee to a squatter settlement organizer who extracts a rent-like payment from them (see Brueckner and Selod (2009) and Marx et al. (2019)). All these rents are

[7]This "ex ante" setting can also be rationalized by considering that the equilibrium captures a steady-state representation of a life-cycle process in which households lose and change jobs (albeit not residential locations) over time.

endogenous and will be determined in equilibrium. Regarding beneficiaries of formal subsidized housing ($h = FS$), We assume that they pay a zero rent to the government. In practice, this simplification contributes to making public housing offers preferable over any other housing option in our equilibrium simulations. This implies that the whole stock of formal subsidized housing is provided to a fraction of income group 1 households (without any vacancy) until other eligible households are rationed out.[8]

The last term on the right-hand side of equation (2) captures construction and maintenance costs of housing structures given expected damages from floods. Here, we treat differently formal private residents and other households. For formal private housing, construction and maintenance costs are borne by developers and thus only indirectly enter the budget constraint of households through housing prices (see the "housing supply" subsection below). Since all other dwellers are assimilated to owner-developers, they bear construction and maintenance costs in the face of flood damages.[9]

Mathematically, as regards maintenance costs, households in formal subsidized housing, informal settlements, and informal backyards ($\mathbb{1}_{\{h \neq FP\}}$) all have to pay a fraction $\rho + \rho_h^{struct}(x)$ of building replacement costs v_h (remember that these housing structures have a fixed size for $h \neq FP$). In addition, dwellers in informal settlements and in informal backyard structures ($h \neq FS$) pay a fraction δ of this amount. This capture the costs of financing construction, i.e. the interest rate paid on a loan contracted to cover construction costs.[10]

Developers We focus on the housing supply per unit of available land. Since informal units cannot be built vertically, it is equal to one for $h = IS$ and $\mu(x)$ for $h = IB$. Moreover, since the supply of (standardized) public housing is exogenous, we abstract from modeling what drives the (political) decision to build houses for $h = FS$. We are left with having to explain the choices of formal private developers.

We assume that formal private developers produce housing with capital and land using a Cobb-Douglas production function with constant returns to scale. Given the very flat gradient of built density in Cape Town (implying an elasticity of substitution between land and capital that would be close to one in a CES specification), we consider that using Cobb-Douglas is a realistic approximation (also see Epple et al. (2010) and Combes et al. (2021)). With this assumption, the produced quantity of formal private housing per unit of land is:

$$s_{FP}(k) = \kappa k^{1-a} \tag{3}$$

where κ is a scale factor, $k = \frac{K}{L_{FP}}$ is the endogenous amount of capital per unit of land available

[8]Note that, due to such rationing, beneficiaries of formal subsidized housing will have greater equilibrium utility than non-beneficiaries within income group 1.

[9]This differs from Monnier et al. (2026) where beneficiaries of formal subsidized housing act as micro-developers and, as such, provide the structures to informal backyarders.

[10]It can also be interpreted as a flow cost of construction since poor households may not have access to financial markets and incrementally build their structures.

for private development, $0 < a < 1$ is the elasticity of formal housing production with respect to land, and $1 - a$ is the capital elasticity.

Formal private developers maximize a profit function (per unit of available land) defined as:

$$\Pi(x,k) = R_{FP}(x)s_{FP}(k) - \left(\rho + \rho_{FP}^{struct}(x) + \delta\right) k - \delta P(x) \tag{4}$$

In equation (4), developers contract a debt over an infinite amount of time to buy unbuilt land from absentee landlords at unit price $P(x)$ and pay an interest δ on this debt in our static setting.[11] They also borrow capital at rate δ (paying a cost δk of borrowed capital per unit of developed land) to produce housing according to equation (3), which they rent out at unit rent $R_{FP}(x)$. Additionally, since they own the structures, they need to cover for both the standard capital depreciation rate ρ, and the expected fraction of capital destroyed due to floods, $\rho_{FP}^{struct}(x)$.

4.2 Housing demand

The size of dwellings is fixed for $h = FS, IS, IB$. For the formal private market, maximizing utility (1) under constraint (2) and plugging back the solution $Q^*(x,i)$ into (1) implicitly defines the optimal housing consumption in location x for income group i conditional on utility u as the solution to:

$$u = \left(\frac{\alpha \tilde{y}_i(x)}{1 + \gamma \rho_{FP}^{content}(x)}\right)^{\alpha} \frac{Q^*(x,i) - q_0}{(Q^*(x,i) - \alpha q_0)^{\alpha}} A(x) \tag{5}$$

We denote this solution $Q^*(x, i \mid u)$. Further imposing that developers do not rent out housing units below a legal threshold $q_{min} > q_0$, the dwelling size for $h = FP$ is thus $Q_{FP}(x,i,u) = max\left[q_{min}, Q^*(x, i \mid u)\right]$.

4.3 Market rent

Plugging back $Q_{FP}(x,i,u)$ into the first-order optimality condition for formal private households (not shown), we obtain the bid rent for a unit of housing type $h = FP$, which corresponds to the maximum amount a household of income-group i would be willing to pay for a unit of formal private housing in location x to attain utility u:

$$\psi_i^{FP}(x,u) = \frac{(1-\alpha)\tilde{y}_i(x)}{Q_{FP}(x,i,u) - \alpha q_0} \tag{6}$$

By doing the same for households living in informal housing, with fixed dwelling size q_I and

[11]In our setting, the price $P(x)$ can be recovered as a function of (endogenous) housing market rent $R_{FP}(x)$ from a zero-profit condition: $P(x) = \frac{\kappa^{\frac{1}{a}} a (1-a)^{\frac{1-a}{a}}}{\delta(\rho+\delta)^{\frac{1-a}{a}}} R_{FP}(x)^{\frac{1}{a}}$.

replacement cost v_I, we obtain the following bid rents for $h = IS, IB$:

$$\psi_i^{IS}(x,u) = \frac{1}{q_I}\left(\tilde{y}_i(x) - \left(\rho + \delta + \rho_{IS}^{struct}(x)\right) v_I - \left(1 + \gamma \rho_{IS}^{content}\right)\left[\frac{u}{(q_I - q_0)^{1-\alpha} A(x) B_{IS}(x)}\right]^{\frac{1}{\alpha}}\right) \tag{7}$$

$$\psi_i^{IB}(x,u) = \frac{1}{q_I}\left(\tilde{y}_i(x) - \left(\rho + \delta + \rho_{IB}^{struct}(x)\right) v_I - \left(1 + \gamma \rho_{IB}^{content}\right)\left[\frac{u}{(q_I - q_0)^{1-\alpha} A(x) B_{IB}(x)}\right]^{\frac{1}{\alpha}}\right) \tag{8}$$

In equilibrium, housing will be allocated to the highest bidding income group in each grid cell where housing of a given type $h = FP, IS, IB$ is available. The upper envelope of bid rents will define the equilibrium market rent $R_h(x)$ across space for $h = FP, IS, IB$. Remember that the rent is considered to be zero for $h = FS$.

4.4 Housing supply

In the formal private sector, profit maximization of developers with respect to capital (per unit of land) in equation (4) defines the optimal quantity of invested capital. Plugging this value back into equation (3) yields:

$$s_{FP}(x) = \kappa^{\frac{1}{a}}\left(\frac{(1-a)R_{FP}(x)}{\rho + \rho_{FP}^{struct}(x) + \delta}\right)^{\frac{1-a}{a}} \tag{9}$$

In the informal backyard sector, utility maximization of (the poorest) households living in formal subsidized housing with respect to the fraction of backyard space that is rented out yields:

$$\mu(x) = \alpha \frac{q_{FS} - q_0}{Y} - (1-\alpha)\frac{\tilde{y}_1(x) - (\rho + \rho_{FS}^{struct}(x))\, h_{FS}}{Y R_{IB}(x)} \tag{10}$$

As mentioned previously, the housing supply per unit of available land is exogenously given for FS housing and is equal to one in (exogenously given) IS plots.

4.5 Population distribution

For the highest bidding income group besides formal subsidized housing recipients (i.e., for $h = FP, IS, IB$), we divide the equilibrium housing supply (per unit of available land) by the optimal dwelling size, and multiply it by the amount of available land to obtain the number of households in each grid cell for each housing type. We obtain:

$$N_i^h(x) = \frac{s_h(x) L_h(x)}{Q_h(x,i,u)} \tag{11}$$

The formula also applies for $h = FS$ with $i = 1$ since only income-group 1 households are eligible for formal subsidized housing.

4.6 Determination of the equilibrium

Following subsections 4.2-4.5, we define an equilibrium as a set $\{u_i, R_h(x), s_h(x), N_i^h(x)\}$ for all i, h, and x where these functions are defined, and where the following constraints hold:

(i) $u_i^h(x) = u_i$ for all $x \in X_{\backslash FS}$ (set of locations occupied by households other than public housing beneficiaries) [*Spatial equilibrium*]. This ensures that there is no profitable deviation in equilibrium: all members of a given income group are indifferent between available location-housing pairs (with the exception of public housing beneficiaries who all have a specific utility level resulting from the exogenous allocation of housing units).

(ii) $P(x) \geq P_A$ for $x \in X_{FP}$ (set of locations occupied by formal private housing dwellers) where P_A is the price of raw agricultural land [*City-edge constraint*]. This ensures that, at the city fringe, absentee landlords are indifferent between selling their land to a developer or engaging in agricultural activities, which endogenously defines city boundary locations.

(iii) $N_i = \sum\limits_h \sum\limits_x N_i^h(x)$ [*Population clearing*]. This ensures that the city hosts all individuals in equilibrium (closed-city resolution).

Note that we opted for a "closed-city" resolution with endogenous utilities as opposed to an "open-city" resolution where utilities are exogenous but population is endogenous (see Fujita (1989)). This is because we do not expect flood awareness policies (the focus of our paper) to generate a large migration response and because closed-city models can easily account for exogenous population change (which could largely reflect the local government's tolerance for informal settlements).

The numerical algorithm we use to solve for the equilibrium implements this definition. Starting from arbitrary utility levels for each income group, with the exception of public housing beneficiaries (condition (i)), we sequentially solve for the quantities determined in subsections 4.2-4.5 and find the corresponding urban extent (condition (ii)). We then compute the error between the predicted populations $\sum\limits_h \sum\limits_x N_i^h(x)$ and the target populations N_i taken from the data (condition (iii)). We then modify utility levels depending on the sign and size of the error terms: the larger the error, the more we increase utility, and conversely (note that $Q_h(x, i, u)$ is increasing in u in equation (11)). The process is iterated until the total absolute error falls below a predefined precision threshold.

Finally, we do not formally prove that the equilibrium exists and is unique. However, the fact that the algorithm converges towards the same results for 250 simulations starting from a wide range of initial values suggests that it does (see **?** for a formal proof in a related setting).

5 Estimation and calibration

5.1 Expected income net of commuting costs

We consider C employment locations, indexed by $c = 1, ..., C$, offering wage w_{ic} to members of income group i. The purpose of this subsection is to estimate the wage w_{ic}, and recover expected income net of commuting costs $\tilde{y}_{ic}$. Note that this quantity is estimated at the benchmark and does not change across counterfactuals. Our model should therefore be interpreted as a spatial equilibrium model of the housing market—without an endogenous labor market—but in which we nest a workplace choice model (see below). In this model, households choose a residential location anticipating they could potentially work in any of the job center but before finding a job in a specific job center.

Workplace choice model Given exogenous employment rates χ_i by income group, households consider expected income $y_{ic} = \chi_i w_{ic}$ in each job center when choosing a residential location.[12] To commute between their residence and potential work locations, households further choose between M potential modes of transportation, denoted by m. For each mode m, residential location x, job center c, and income group i, households of type j face commuting cost:

$$t_{mj}(x, c, w_{ic}) = \underbrace{\chi_i \left(\tau_m(x, c) + \delta_m(x, c) w_{ic}\right)}_{\overline{t_m}(x,c,w_{ic})} + \epsilon_{mxcij}$$

where $\tau_m(x, c)$ is the monetary transport cost and $\delta_m(x, c)$ is the opportunity cost associated with the fraction of time spent commuting, such that $\overline{t_m}(x, c, w_{ic})$ measures the total expected deterministic cost of commuting. ϵ_{mxcij} is a random component which follows a Gumbel minimum distribution of mean 0 and scale parameter $\frac{1}{\lambda}$. It accounts for heterogeneity in individual preferences for each origin-destination commuting pair and transport mode (see Teulings et al. (2018) for a similar approach). It is unobservable to households when making location choices.

Commuters pick the transport mode that minimizes their commuting cost. Due to the properties of the Gumbel minimum distribution (extreme value theory), this comes down to:

$$\underset{m}{min} \left(t_{mj}(x, c, w_{ic})\right) = -\frac{1}{\lambda} log \left(\sum_{m=1}^{M} exp\left[-\lambda \overline{t_m}(x, c, w_{ic})\right] \right) + \eta_{xcij}$$

where η_{xcij} also follows a Gumbel minimum distribution with same underlying parameters.

Note that households still have idiosyncratic preferences for residential-work location pairs once optimal modal choice is accounted for.

For a given residential location x, households therefore choose a workplace c that maximizes their

[12] We treat households as individual workers, which allows us not to consider the more complex issue of joint household location decisions in the case of non-single-member households.

expected income net of commuting costs, solving the program: $\underset{c}{max}\left[y_{ic} - \underset{m}{min}\left(t_{mj}(x,c,w_{ic})\right)\right]$. Since this corresponds to the latent variable formulation of a multinomial logit model (discrete choice theory), the expected income net of commuting costs $\tilde{y}_i(x) \equiv \mathbb{E}\left[y_{ic} - \underset{m}{min}\left(t_{mj}(x,c,w_{ic})\right)\right]$ that enters the main housing choice model can be defined as:

$$\tilde{y}_i(x) = \sum_{c=1}^{C} \pi_{c|ix} \left[y_{ic} + \frac{1}{\lambda} \log \left(\sum_{m=1}^{M} \exp\left[-\lambda \overline{t_m}(x,c,w_{ic})\right] \right) \right]$$

with:

$$\pi_{c|ix} = \frac{\exp\left[\lambda y_{ic} + \log\left(\sum\limits_{m=1}^{M} \exp\left[-\lambda \overline{t_m}(x,c,w_{ic})\right]\right)\right]}{\sum_{k=1}^{C} \exp\left[\lambda y_{ik} + \log\left(\sum\limits_{m=1}^{M} \exp\left[-\lambda \overline{t_m}(x,c,w_{ic})\right]\right)\right]}$$

Commuting cost calibration Based on the 2011 National Census, we consider employment rates $\chi_{i=1,\dots,4} = [0.57, 0.97, 0.96, 0.97]$.

We also assume that the monetary transport cost takes functional form $\tau_m(x,c) = T(F_m + d_{xc}V_m)$, where F_m and V_m are mode-specific fixed and variable cost components, and d_{xc} is the distance (in km) between points x and c (from the City of Cape Town's Transport Model). In our simulations, we consider $C = 185$ main job centers. T is a constant equal to the average number of trips per year per household in full employment, that we set to 470.

We consider five transportation modes that are prevalent in Cape Town: train, bus, minibus/taxi, car, and walking. Walking is free. For the three public transportation modes we estimate the fixed and variable cost components by regressing transport cost statistics on residence-workplace distances taken from Roux (2013). Considering estimates for the average price of a car, and that materials roughly account for half of this value, we apply capital depreciation rate ρ to this cost component to obtain the yearly vehicle fixed cost, that we divide by T. For the vehicle variable cost, we consider fuel prices and energy efficiency data from the International Energy Agency to obtain a price per km at the benchmark year. Table 1 summarizes the values we find for each mode:

Table 1: Summary table of monetary transport cost parameters

Mode	F_m	V_m
Train	4.48	0.164
Bus	4.32	0.785
Minibus/taxi	6.24	0.522
Car	10.00	1.157

Note: OLS regression results.

As for the opportunity cost fraction of time spent commuting, $\delta_m(x,c)$, we assume that it is equal to the fraction of working time spent commuting, or that the opportunity cost of not working is equal to one. Here, we consider the labour supply elasticity as perfectly inelastic

(8 hours a day), and rule out the possibility that workers may commute out of their leisure time.[13] The time spent commuting is directly given by the City of Cape Town's Transport Model, except for walking, which we assume to have a 4km/h speed.

Estimation of λ and w_{ic} From the workplace choice model, we can derive the expected number of residents of income group i choosing to work in c, denoted W_{ic}, provided that we know the number of residents of income group i with their residence in x, denoted $N_i(x)$, in all locations x. This yields:

$$W_{ic} = \chi_i \sum_x \pi_{c|ix} N_i(x)$$

Note that summing this relation over job centers c yields a labor market clearing condition.

Since the values for W_{ic} and $N_i(x)$ are respectively given by the City of Cape Town's Transport Model and the National Census, we first solve the above equation for w_{ic}, separately for discrete values of parameters λ. This is done numerically starting with average values $\overline{w_i}$ for w_{ic} (captured in $\pi_{c|ix}$) on the right-hand side, and updating them iteratively until the target population is reached on the left-hand side. We then pick the $\{\lambda, w_{ic}\}$ pair that best fits the aggregate distribution of residence-workplace distances from Cape Town's Transport Survey: we select the pair identified by $\lambda = 13.96$.

Going back to the workplace choice model, we can now calculate the value of expected income net of commuting costs $\tilde{y}_i(x)$.

5.2 Housing production function parameters

By plugging the value of formal private housing supply (per unit of available land) from equation (11) into the left-hand side of equation (9), and considering that $R_{FP}(x) = \frac{(\delta P(x))^a (\rho+\delta)^{1-a}}{\kappa a^a (1-a)^{1-a}}$ (zero-profit condition), we obtain the following log-linear relation, that we estimate at the "sub-place" (census tract) level s using the property transaction data:

$$log(N_s^{FP}) = \gamma_1 + \gamma_2 log(P_s) + \gamma_3 log(Q_s) + \gamma_4 log(L_s^{FP}) + \epsilon_s$$

where $\gamma_1 = log\left(\kappa(\frac{1-a}{a})^{1-a}\right)$, $\gamma_2 = 1 - a$, $\gamma_3 = -1$, and $\gamma_4 = 1$.

As this theoretical relation only holds for urban formal private housing, we exclude sub-places in the bottom quintile of property prices where more than 5% of households are reported to live in informal housing, or which we classify as rural (i.e., large areas in the highest surface quintile, where less than 60% of the land can be developed, or located more than 40km away from the CBD). We also exclude the poorest income group who may benefit from public housing options. Table 2 shows the results.

The fact that the estimated values for $\widehat{\gamma_3}$ and $\widehat{\gamma_4}$ are reasonably close to -1 and 1, respectively,

[13] This is not a strong assumption: in South African cities, workers commuting from far away reduce their working hours so as not to travel by night in order to reduce their exposure to crime.

Table 2: Estimated coefficients from log-linear regression on property transaction data

$\widehat{\gamma}_1$	$\widehat{\gamma}_2$	$\widehat{\gamma}_3$	$\widehat{\gamma}_4$
-3.763	0.242	-0.895	0.995
(0.981)	(0.069)	(0.094)	(0.070)

Note: OLS regression results. Standard errors in parentheses.

suggests that the relation is indeed well identified. Solving for a and κ with the estimated values for $\widehat{\gamma}_2$ and $\widehat{\gamma}_1$, respectively, yields $a = 0.758$ and $\kappa = 0.031$.

5.3 Utility function parameters

Stone-Geary specification We set basic need in housing $q_0 = 4m^2$ to reflect the very minimum size of informal units from Rice et al. (2023). Then, we set the income elasticity of housing expenditure $1 - \alpha = 0.25$ following Finlay and Williams (2025), who estimate this term for more general non-homothetic preferences.

Amenity index As for the estimation of the amenity index $A(x)$, we leverage property price data for formal private housing and information on exogenous amenities from the City of Cape Town, defined at the sub-place (census tract) level s. Let us define the observed amenity index as:

$$A_s = \prod_n (a_{n,s})^{\vartheta_n} \epsilon_{A,s}$$

where $a_{n,s}$ are exogenous amenity dummies from the data, ϑ_n are their respective elasticities, and $\epsilon_{A,s}$ is an error term.

Taking housing rents from the data and inverting equation (1) for formal private dwellers (essentially the two richest income groups, i.e. income groups 3 and 4) also yields the amenity index A_s as a function of the dominant income group's utility level $u_{i(s)}$. We can therefore estimate the log-linearization of the above equation for discrete values of $u_{i(s)}$. We pick the value set that minimizes the error $\epsilon_{A,s}$, and use the estimated ϑ_n to recover the predicted value of $A(x) = \prod_n (a_{n,x})^{\vartheta_n}$ at the grid cell level. Table 3 shows the corresponding regression results.

Informal disamenity index The disutility parameters for informal housing types (B_h, $h = IS, IB$) are recovered as benchmark model residuals to match the total population living in informal housing units that is observed in the data.[14] We find $B_{IS} = 0.78$ and $B_{IB} = 0.80$.

Table 4 summarizes the values and identification sources of all estimated parameters.

[14]This is akin to recovering amenity terms from model inversion in QSMs.

Table 3: Results of the log regression of predicted amenity index on exogenous amenities

	$log(A_s)$
Prox. to district park (<1km)	-0.051
	(0.043)
Prox. to ocean (<2km)	0.063
	(0.041)
Prox. to ocean (<4km)	0.011
	(0.046)
Prox. to urban heritage site (<2km)	0.185
	(0.052)
Within airport noise cone	-0.007
	(0.067)
Slope (between 1 and 5%)	0.144
	(0.034)
Slope (>5%)	0.121
	(0.049)
Prox. to biosphere reserve (<2km)	0.098
	(0.038)
Prox. to train station (<2km)	0.039
	(0.035)
Constant	-0.458
	(0.038)
Obs.	307
R^2	0.160
F-stat	6.268
D.f.	297

Note: OLS regression results. Standard errors in parentheses.

Table 4: Estimated parameters

Parameter	Value	Source of identification
a	0.758	Predicted population according to housing characteristics in the formal private sector
κ	0.031	Predicted population according to housing characteristics in the formal private sector
λ	13.96	Distribution of commuting pairs and distances
B_{IS}	0.78	Spatial distribution of informal settlements (model inversion)
B_{IB}	0.80	Spatial distribution of basic backyard structures (model inversion)

Note: The table presents the exogenous parameters that are estimated based on the model predictions and input data, along with their value and the moments used to match the predictions with the data.

5.4 Externally calibrated parameters

Table 5 summarizes the values and data sources for additional model parameters that are externally calibrated.

Table 5: Chosen parameters

Parameter	Value	Data source
q_{FS}	110 (m^2)	HRSC Backyarding Report (2019)
Y	70 (m^2)	HRSC Backyarding Report (2019)
q_I	14 (m^2)	HRSC Backyarding Report (2019)
v_I	3,000 (ZAR)	HRSC Backyarding Report (2019)
v_{FS}	127,000 (ZAR)	HRSC Backyarding Report (2019)
γ	0.27	Quantec (HH budget breakdown)
ρ	0.025	Viguié et al. (2014)
δ	0.043	World Development Indicator database (2016)
P_A	807.2 (ZAR/m^2)	Property transaction data (CoCT)
β	0.25	Finlay and Williams (2025)
q_0	4 (m^2)	Rice et al. (2023)
q_{min}	31.6 (m^2)	Dwelling size data (CoCT)

Note: The table presents the exogenous parameters that are directly taken from external sources, along with their value and their data source.

q_{min} is a legal requirement, whereas q_{FS}, q_I and Y are estimates from the field. The fact that they are fixed values reflects the relative standardization of public housing schemes and of informal construction technologies. Although these are simplifications, we consider them as good first-order approximations.

The replacement values of formal subsidized, v_{FS}, and informal buildings, v_I (both in informal settlements and in the backyards of formal subsidized housing), are based on expert assessments of material costs, along with surveys of construction technologies used in the field. The share of composite good that is vulnerable to floods, γ, corresponds to the average budget share households spend on durable and semi-durable goods. These goods are not consumed immediately and are kept within the house.

The agricultural land price, P_A, corresponds to the highest decile in the property price data when selecting rural areas only. The capital depreciation rate, ρ, is 2.5 times higher than the one used by Viguié et al. (2014) for Paris, which accounts for the overall lower construction quality in Cape Town. Finally, the interest rate, δ, corresponds to the 4-year average up to the benchmark year (2011) in the World Development Indicator database, so as to smooth its volatility.

6 Simulations

In what follows, we run our benchmark model and conduct two scenario comparisons. In the first one, we compare our benchmark to a version of the model where agents do not anticipate flood risks. In the second simulation comparison, we compare our benchmark (with perfect information on flood risks) to a version of the model where flood risks are updated to account for climate change.

Observe that when households do not anticipate flood damages, equilibrium utilities are defined irrespectively of the realization of the likelihood and intensity of flood damages. This means that comparing equilibrium utilities with and without anticipation would be misleading. Instead, we compare expected damage values across simulations, both in absolute terms (to calculate aggregate economic losses) and in relative income terms (to proxy for individual welfare losses).

6.1 The role of information

Aggregate damages In Figure 4, we represent separately for each flood type the distribution of expected aggregate annual monetary damages across housing and damage types (in red for damages to structures and in blue for damages to contents) for our two simulations (in dark color for the equilibrium with anticipation, and in light color for the equilibrium without anticipation).

Several comments are in order. First, damages for formal housing are consistently larger than for informal housing, although the latter is relatively more vulnerable to flooding (due to the building materials used and the locations of informal settlements in flood-prone zones). This is because the replacement cost of informal settlements and backyard structures is based on construction costs that are fairly low. Formal private housing, on the other hand, is built according to higher and more costly standards, and will typically use land more intensively. The replacement cost of formal subsidized housing structures lies somewhere in between. In addition,the poorest income group is bid away from the formal private segment in practice. The higher amount of income available for consumption in richer income groups hence explains the relative importance of content damages in this category.

Second, as expected, damages are higher in the no-anticipation case than in the perfect-anticipation case. Anticipation also seems to play a relatively more important role in avoiding

Figure 4: Aggregate flood damage estimates across flood types
with (w/) and without (w/o/) anticipation of flood risks

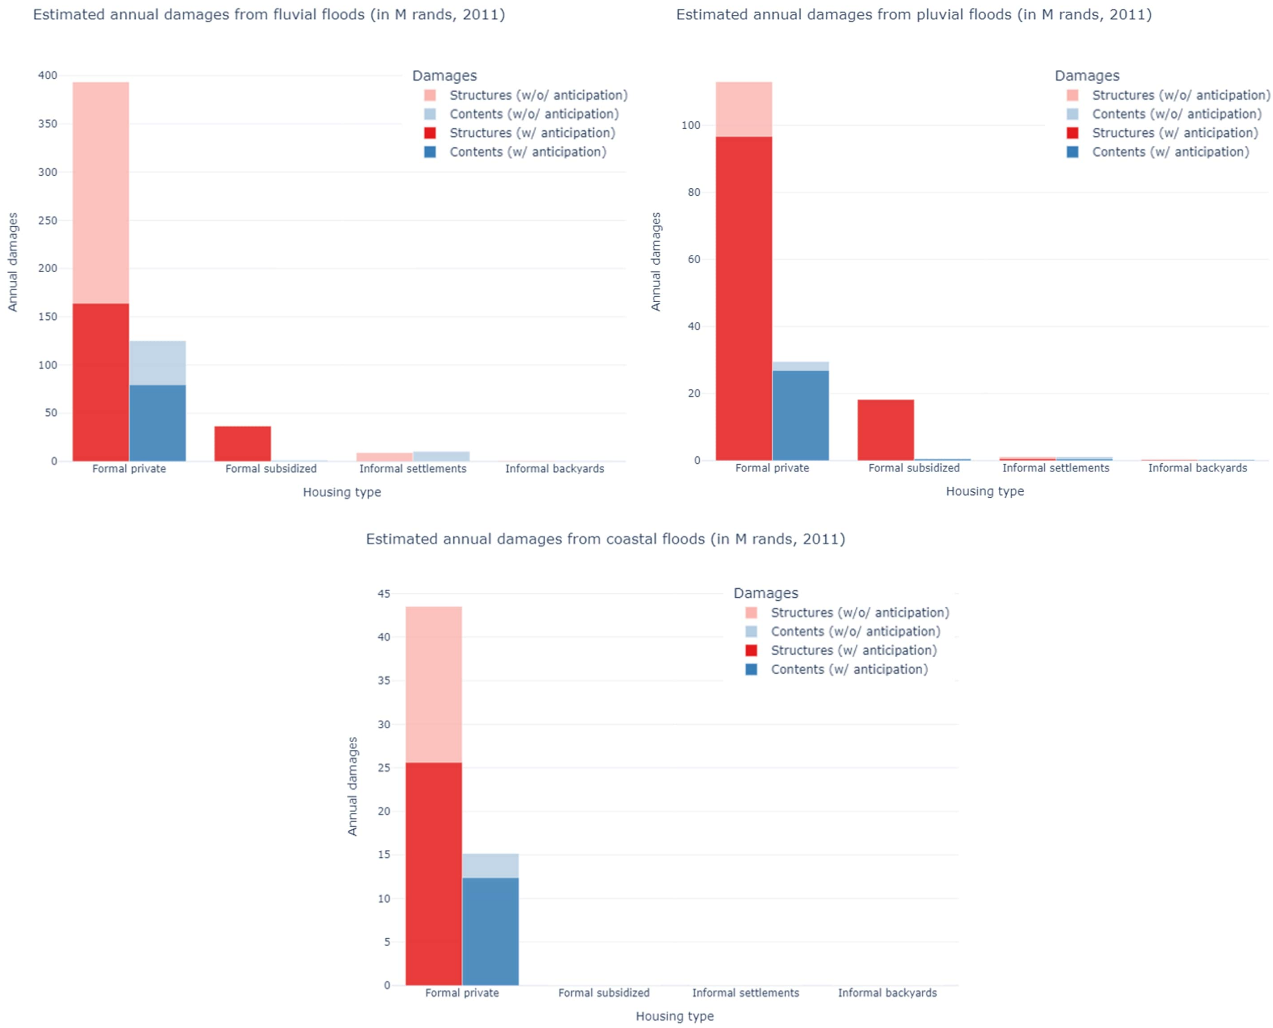

Source: Authors' simulations. Red is for damages to structures, blue for damages to contents, dark colors for the equilibrium with anticipation, and light colors for the equilibrium without anticipation. Dark color bars are overlaid over light color bars.

damages in the formal private housing sector. This is because it is the least constrained of the housing submarkets, either in terms of available locations, or given the less stringent restrictions on supply. Moreover, anticipation plays a contrasted role across flood types. This can be interpreted in terms of existing constraints: pluvial flood zones are typically more dispersed than fluvial or coastal flood zones, and are hence harder to avoid through relocation. As for coastal flood zones, they are typically more attractive, hence a higher willingness of households to stay in these locations and absorb the added expected cost, all the more so for rich households with a low marginal utility of income.

Third, damages also significantly vary in magnitude across flood types. Table 6 summarizes aggregate damage estimates per flood type in our two simulations, with values converted from 2011 ZAR to 2021 USD using an inflation rate over the period of 64% and an exchange rate of 14.78 ZAR/USD. Note that total damages are smaller than the sum of individual categories as we avoid double counts when flood zones overlap.

Table 6: Summary table of aggregate damage estimates w/ and w/o/ anticipation

	Anticipation	No anticipation
Fluvial	280M (ZAR, 2011)	570M (ZAR, 2011)
	31M (USD, 2021)	*63M (USD, 2021)*
Pluvial	145M (ZAR, 2011)	165M (ZAR, 2011)
	16M (USD, 2021)	*18M (USD, 2021)*
Coastal	50M (ZAR, 2011)	75M (ZAR, 2011)
	6M (USD, 2021)	*8M (USD, 2021)*
Total	448M (ZAR, 2011)	750M (ZAR, 2011)
	49M (USD, 2021)	*83M (USD, 2021)*

Source: Authors' simulations. Each row gives the expected annual damages estimated for each flood type, with and without flood risk anticipation from agents in the model. The total estimates only account for the most severe flood type when flood zones overlap, so as to avoid double counting. South African rands (ZAR, 2011) are converted to US dollars (USD, 2021) using an inflation rate of 64% from 2011 to 2021 and an exchange rate of 14.78 ZAR/USD in 2021.

For comparison, Gonzales et al. (2023) estimate aggregate annual flood damages in Cape Town at around 16M USD (at 2021 values). The key difference between our results lies in the way the replacement cost of vulnerable assets is determined. In their paper, a standard capital building stock per unit of land is calibrated and yields flood damages based on depth-damage conversions. In our paper, the capital used in buildings is an endogenous outcome that can quickly rise in attractive areas (where land is used intensively), and we also take housing contents into consideration.[15]

At any rate, our paper shows that material flood damages are high (potentially higher than previously thought), and that information is a key policy lever to reduce them, with a reduction potential of up to 41% when looking at the difference between total damages in our two simulations.

Absolute damage distribution Figure 5 shows how flood damages spread out across the city in the no-anticipation case, and the absolute changes in damages compared with the benchmark anticipation case. The largest damage values are concentrated in populated areas along rivers near the city center and to the north, where large job centers are located. These are also areas where damage values change the most compared with the benchmark case, suggesting an important adaptation of households following an information shock.

In relative damage terms (calculated as an equivalent share of expected income net of commuting costs, henceforth "net income"), values would appear to be higher in the north than in the city center, and even more to the east, where poorer households live. As households of different types move across locations in the two simulations, it is not possible to properly estimate the changes

[15]In fact, our total estimate for the capital used in the whole city is in line with that of the authors. The reason we find higher damages therefore boils down to a composition effect, as we are able to simulate how capital gets distributed across space, notably in flood-prone zones.

Figure 5: Spatial absolute flood damage estimates w/ and w/o/ anticipation (ZAR, 2011)

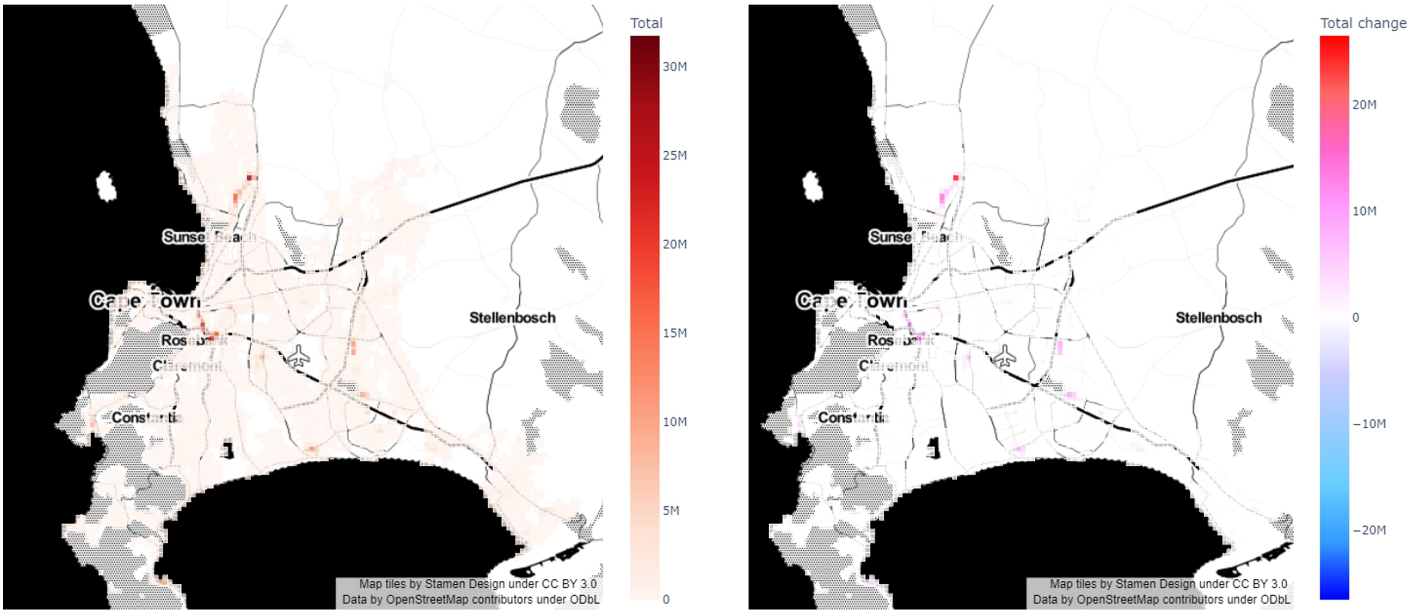

Note: The map on the left shows the values under no flood risk anticipation. The map on the right shows the change (in absolute terms) it represents compared with the benchmark (perfect anticipation).

in relative damages at the grid-cell level. Instead, we now turn to their aggregate distribution per income group to proxy how flood damages translate into welfare losses.

Relative damage distribution Figure 6 represents the aggregate distribution of relative flood damages per income group. Since there is a high mass of households who do not live in flood-prone areas, the graphs only consider the subpopulation living in areas with a positive risk of floods. In the interest of space, we only show results for fluvial flood risks (since they are the most impactful in households' adaptive behaviour).[16]

We see that richer households are less likely to live in flood-prone zones, even in the no-anticipation case. This is because these zones typically are low-amenity areas located far from job centers. Interestingly, conditional on flood exposure, the tail of relative damage distributions does not necessarily become thinner as we consider richer households, even in the anticipation case. This is because, as marginal utility decreases with income, richer households become comparatively less sensitive to flood damages.

As a side note, the long tail of relative damage values for the poorest income group (as seen on the horizontal axis) does not only reflect their low income levels, but also to the fact that replacement costs of formal subsidized housing are very high (compared with informal housing). In the model, we make the assumption that public housing providers would cover extreme destruction costs in case they become unaffordable to the beneficiaries.

To visualize more clearly the adaptation response, Table 7 summarizes average relative damage estimates per income group. The damage reduction generated by the information shock is

[16]Other results are available upon request.

Figure 6: Relative fluvial flood damage distribution by income group w/ and w/o/ anticipation

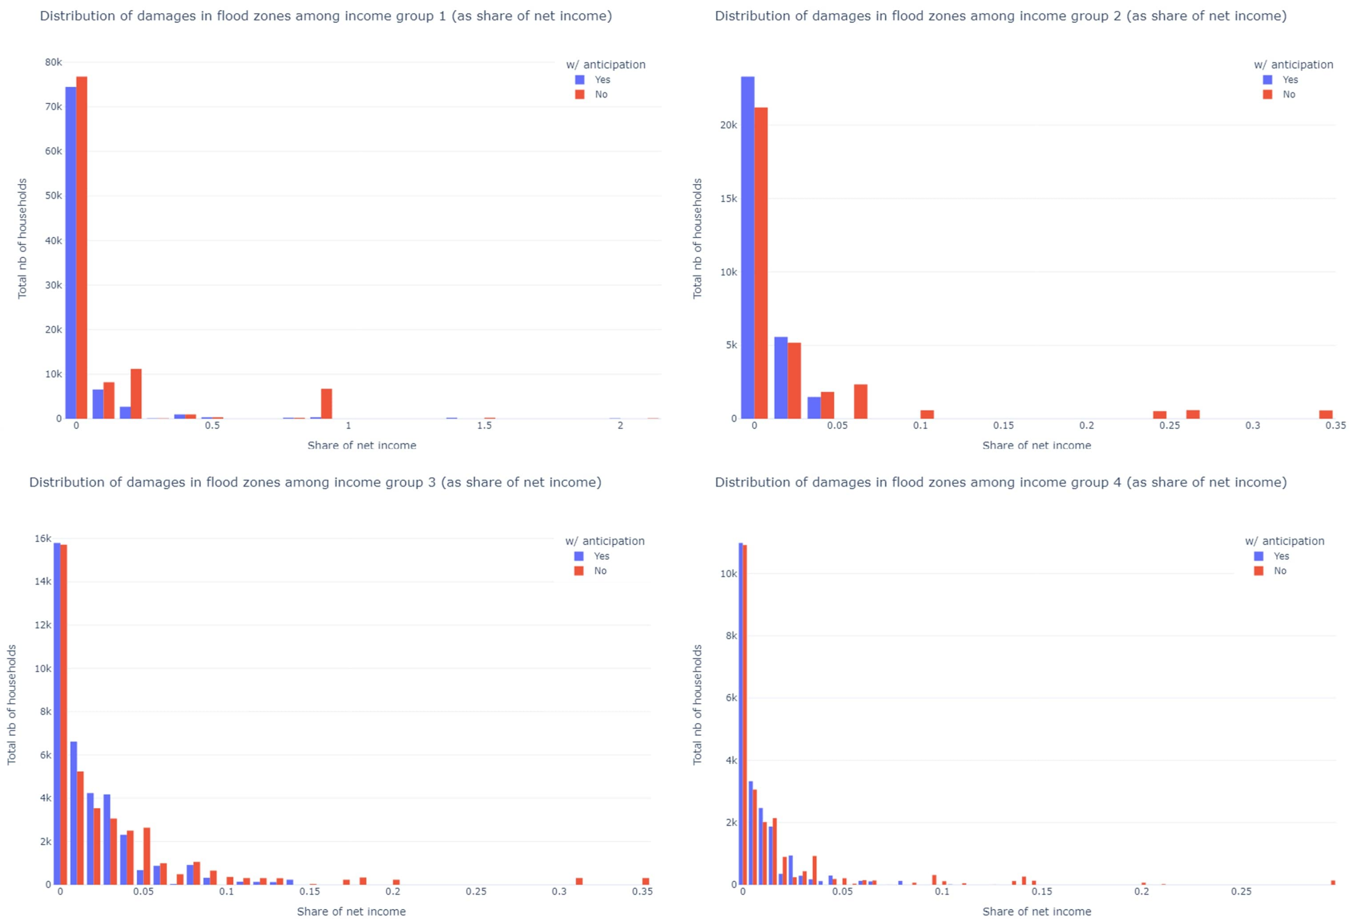

Source: Authors' simulations. Each panel shows the distribution of expected fluvial flood damages across the population as a share of equivalent income net of commuting costs, separately for each income group. Blue bars stand for simulation results with flood risk anticipation, and red bars for simulation results without flood risk anticipation.

equivalent to more than 4% of net income for 20% of the total population initially exposed to flood risks. This is thanks to households relocating to areas with lower flood risks (potentially leaving flood-prone zones entirely) and readjustments in the housing supply.

Looking at heterogeneity by income group, we find that the poorest income group makes up roughly half of the total exposed population. Households in this group are also more likely to live in flood-prone areas and face the largest expected damages relative to their net income. They therefore tend to gain the most from having information, with total income equivalent gains rising to 7% for this group. As we saw, this is in spite of their choice set being more constrained. This leaves room for even further gains, for instance under more accommodating housing policies.

Flood damages and policy responses therefore appear to be substantive, not only in absolute but also in relative terms, even though the household and housing types that are the most affected are not the same for absolute or relative damages. We now give more details on how adaptation takes place in practice.

Table 7: Summary table of relative damage estimates w/ and w/o/ anticipation

	Anticipation	No anticipation
Inc. group 1		
Inc. equiv. damages	3.37%	9.83%
Exposed pop.	87,582	104,723
Group share	22%	26%
Inc. equiv. gain	*+7.01%*	
Inc. group 2		
Inc. equiv. damages	0.60%	2.57%
Exposed pop.	30,363	32,789
Group share	18%	19%
Inc. equiv. gain	*+2.01%*	
Inc. group 3		
Inc. equiv. damages	1.79%	3.08%
Exposed pop.	36,593	38,620
Group share	12%	13%
Inc. equiv. gain	*+1.38%*	
Inc. group 4		
Inc. equiv. damages	0.76%	1.59%
Exposed pop.	21,192	22,624
Group share	13%	14%
Inc. equiv. gain	*+0.88%*	
All		
Inc. equiv. damages	2.25%	6.39%
Exposed pop.	175,730	198,656
Group share	17%	19%
Inc. equiv. gain	*+4.40%*	

Source: Authors' simulations. Each panel gives, separately for each income group, expected flood damages as a share of equivalent net income, the corresponding exposed population and the group share it represents, across simulations with and without flood risk anticipation. An additional row gives the income equivalent gain obtained from damage reduction when introducing anticipation.

Population relocation Figure 7 shows the spatial population distribution in the non-anticipation case, and the corresponding changes (in absolute terms) compared with the benchmark anticipation case. As already mentioned, the population density is highest in informal settlement areas in the east (which exist alongside formal private housing units) and where public housing (hence informal backyard) units tend to also be located. Interestingly, these are the areas where there appears to be the most moves between the two simulations: poor households appear to be relatively mobile, even though they do not seem to move far from their original locations.

Figure 7: Spatial population distribution w/ and w/o/ anticipation

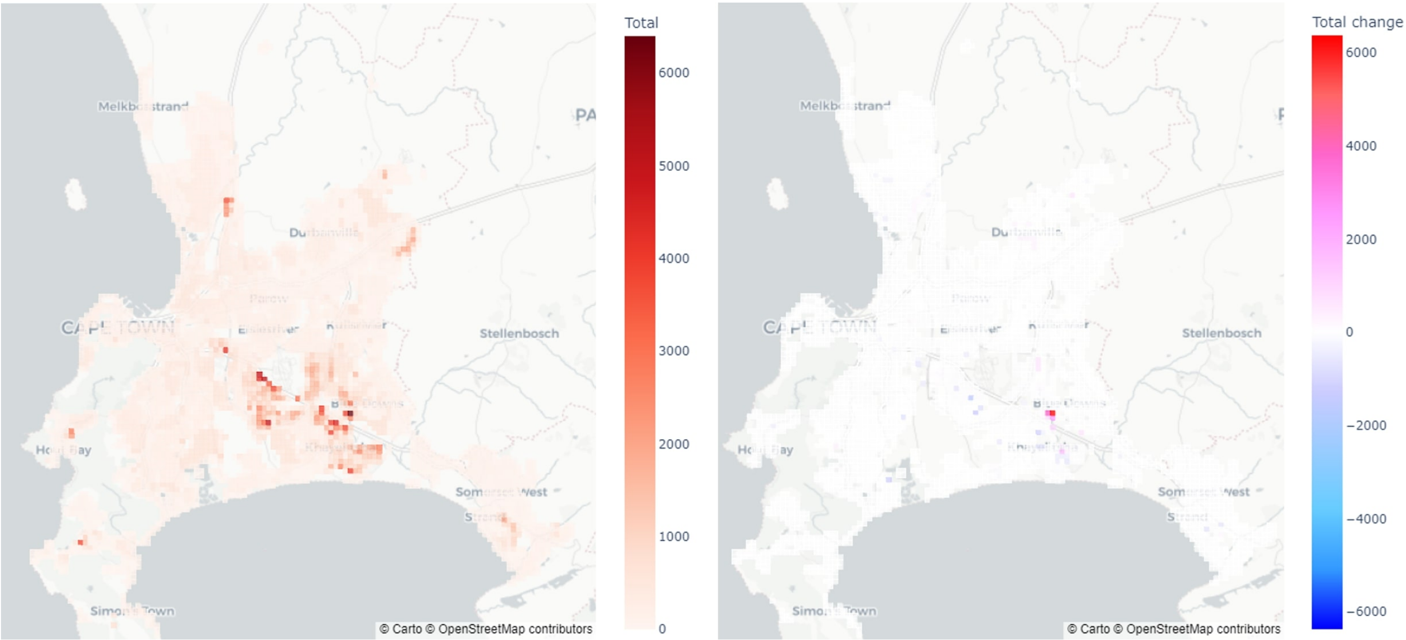

Note: The map on the left shows the values under no flood risk anticipation. The map on the right shows the change (in absolute terms) it represents compared with the benchmark (perfect anticipation).

Changes in rents Figure 8 shows spatial changes in rents (in absolute terms) across the three endogenous housing submarkets in the model. As could be expected, rents generally increase with population as both variables capture the willingness to pay of households across the two simulations. Interestingly, they also change in places where there are no strong population changes, and even more so where we expect relatively poor households to live. This suggests that adaptation also occurs with households demanding lower rents to be compensated for the cost of living in flood-prone zones (or willing to pay more not to be bid away from safe places). This effect is stronger for vulnerable populations who are more sensitive to changes in income. Incidentally, falling rents can also reflect a shrinking housing supply that constrains households to consume less housing. This is the case in central valuable areas where flood damages are found to decrease.

All in all, poor households seem to be reactive to the information shock, with a relocation response that allows them to reduce their exposure to flood risks, and a market response that allows them to be partly compensated through changes in rents when they are able or willing to

Figure 8: Absolute change in spatial rent (rands/m²) per housing type in the no-anticipation case (compared with perfect anticipation)

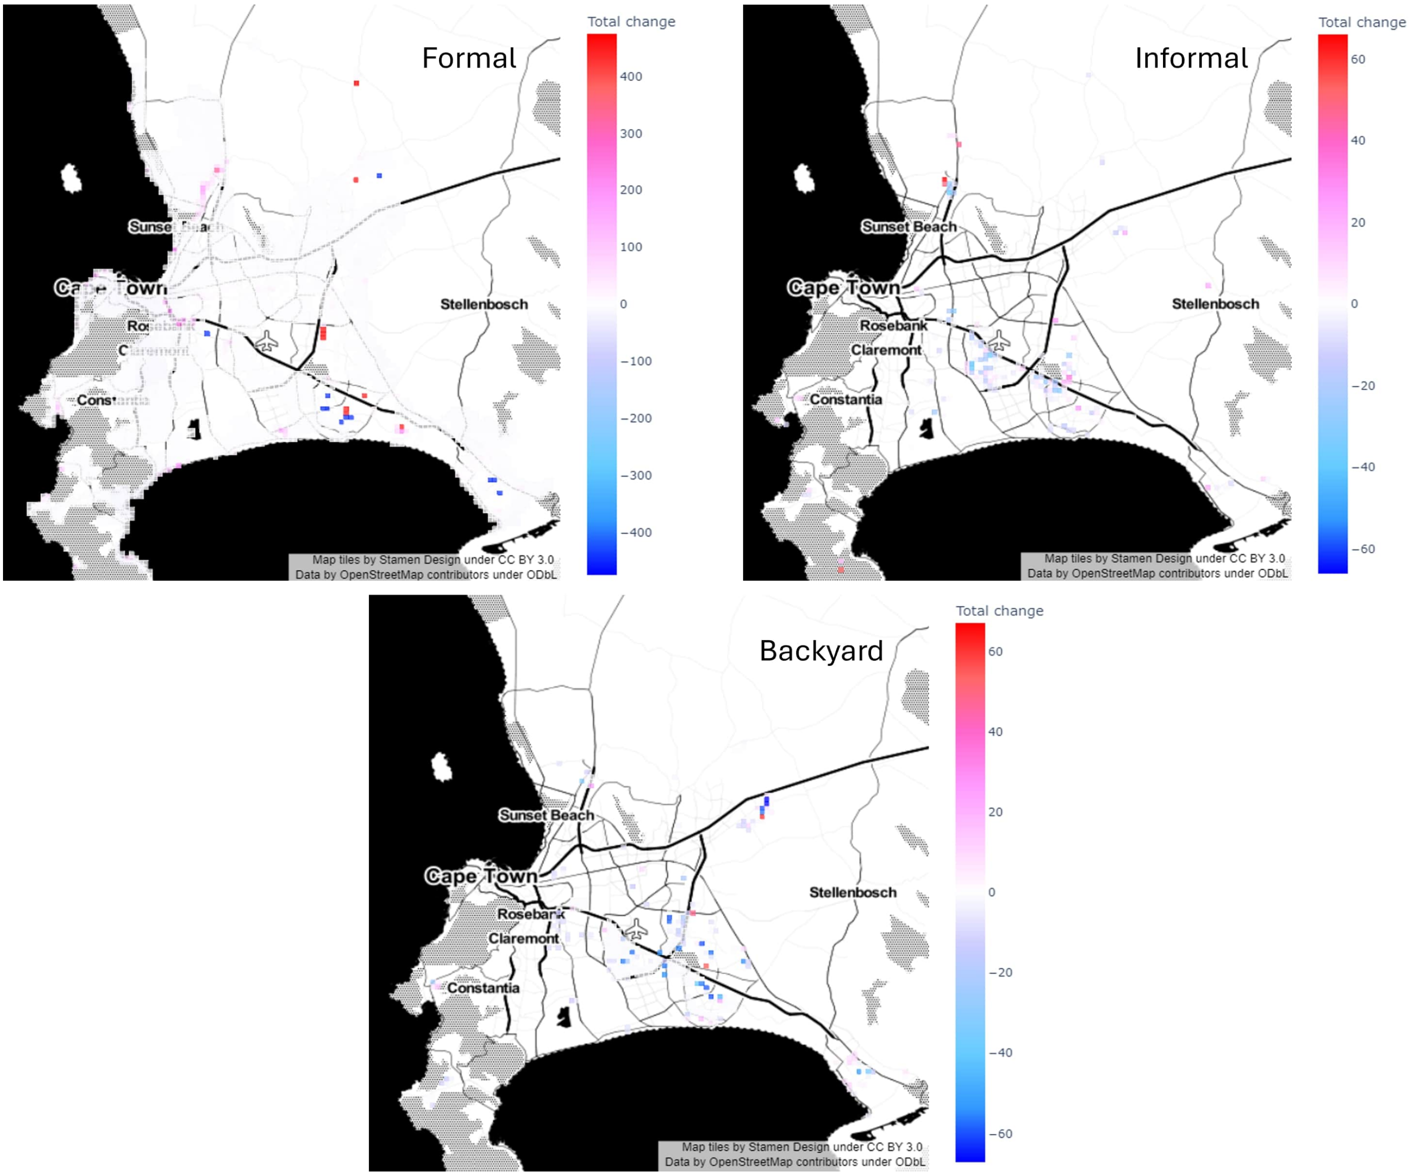

Note: Each map represents by how many rands/m² the market rent would increase (in red) or decrease (in blue) in our simulation without flood risk anticipation (compared with perfect anticipation), separately for each housing type.

move away.

We now turn to our simulations of climate change adaptation.

6.2 Climate change

Aggregate damages As in the previous comparison, we start by representing in Figure 9 aggregate damages per flood type across our simulations, this time with and without climate change (but with flood risk anticipation in both instances). The largest change now occurs for pluvial flood risks, and not fluvial flood risks. The explanation is the same as before: since pluvial flood zones are more spread out, increased risks in these zones are harder to avoid. Here again, coastal flood zones fall in between as they tend to be high-amenity areas where (rich) households prefer to absorb added costs rather than changing location. Otherwise, the pattern is similar to the previous simulation comparison, with substantial increases in flood damages when introducing climate change.

Figure 9: Aggregate flood damage estimates across flood types w/ and w/o/ climate change

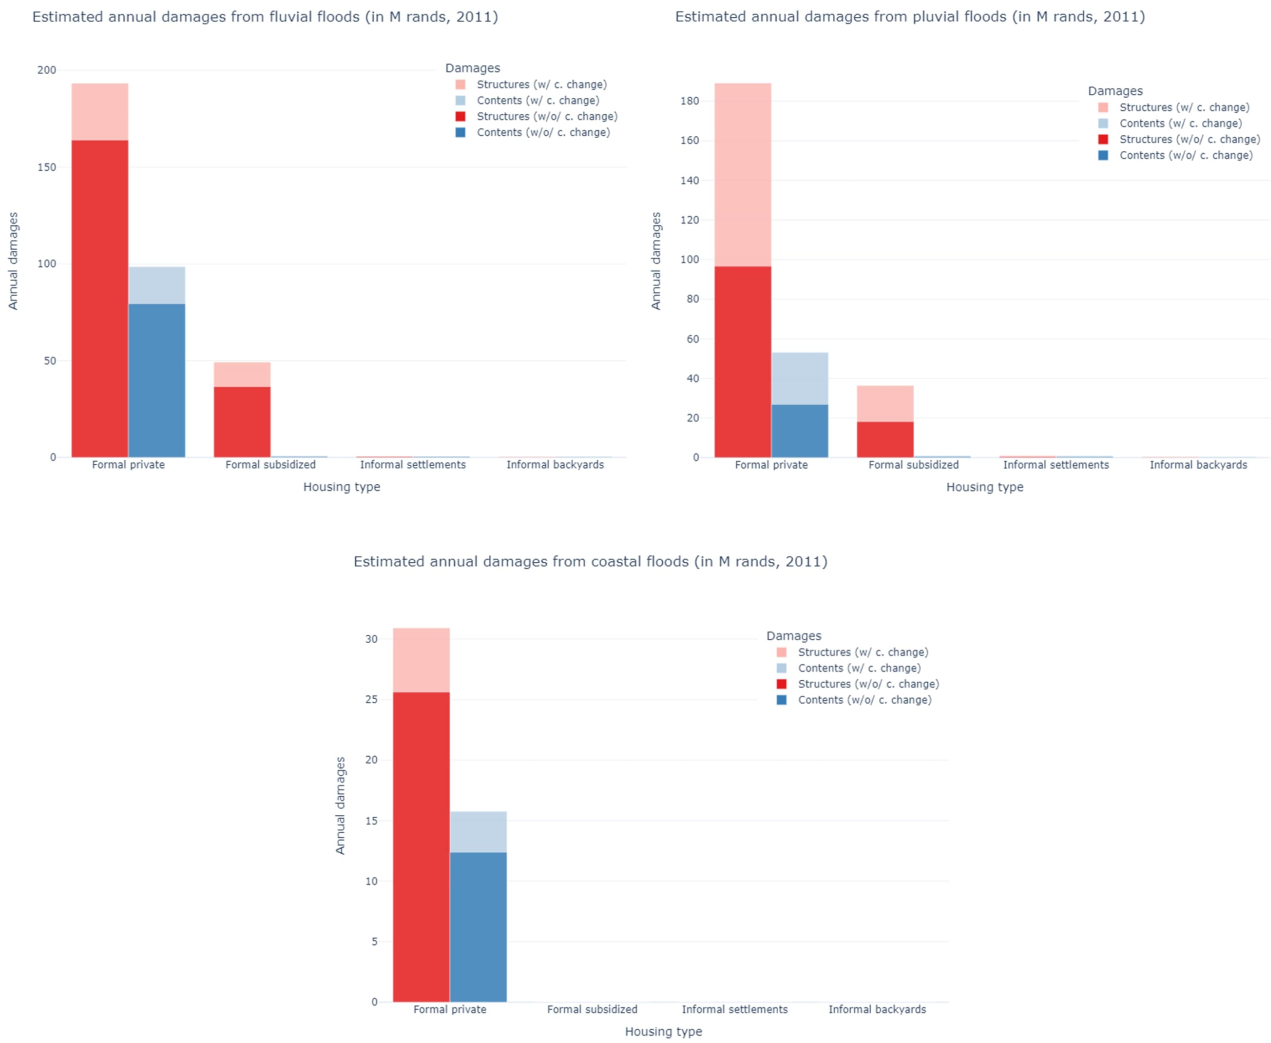

Source: Authors' simulations. Red is for damages to structures, blue for damages to contents, dark colors for the equilibrium without climate change, and light colors for the equilibrium with climate change. Dark color bars are overlaid over light color bars.

We do not reproduce here the aggregate damage estimate table whose primary function was to compare our results with those of Gonzales et al. (2023), who do not consider climate change in their analysis.

Absolute damage distribution Figure 10 shows the spatial distribution of absolute damages in our benchmark case and their increase when taking climate change into account. We see that the change across scenarios is more widely spread (but also smaller per unit of land) than in the previous comparison (as pluvial flood risks are more spread out, but also less severe than fluvial ones).

Figure 10: Spatial absolute flood damage estimates w/ and w/o/ climate change (ZAR, 2011)

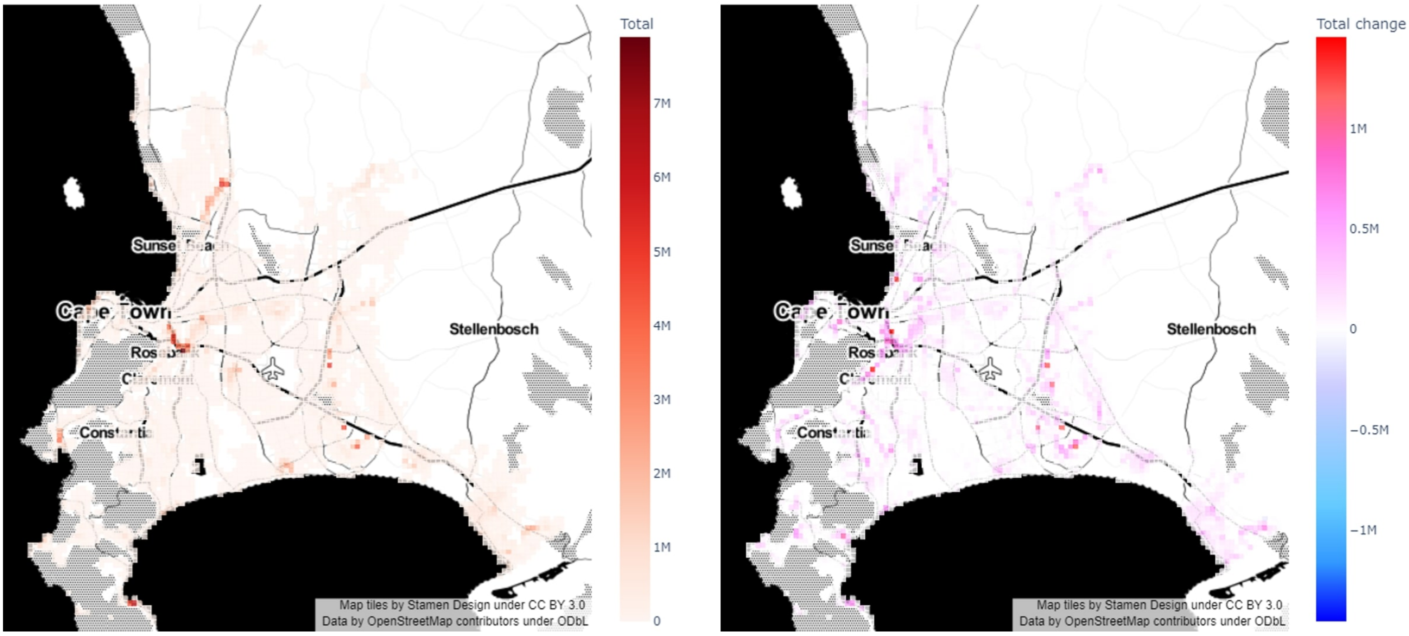

Note: The map on the left shows the values at the benchmark (perfect anticipation). The map on the right shows the change (in absolute terms) when considering climate change.

Relative damage distribution The relative damage distribution per income group is shown in Figure 11 for pluvial flood risks, since they are now the most impactful flood type across the two scenarios. Compared with fluvial flood damages shown in the previous section, there are now more exposed households, and the distribution tails are typically thinner and shorter. As expected, we see exposure shifting from low-damage brackets to high-damage brackets when introducing climate change.

Table 8 summarizes the results. As before, the poorest income group is the most affected, and this time tends to lose the most from the change. Overall, we estimate the relative damages from climate change to roughly 0.3% of the equivalent net income for 90% of the total population. Assuming the same exposed population benchmark as in the previous comparison, this would correspond to a loss of 1.7%, which is less important in absolute terms but not negligible.

Households therefore appear to be less able or willing to adapt to the consequences of climate change regarding pluvial, as opposed to fluvial flood risks. Let us now turn to population relocation to assess spatial adaptation strategies.

Population relocation Figure 12 shows the spatial population distribution at the benchmark, and how it evolves following the climate change shock. The pattern is similar to that of the first simulation comparison as households are avoiding risks in fluvial flood zones.

Changes in rents Figure 13 represents the spatial changes in rents across the three housing submarkets. In this case as well, there seems to be adjustments even when households do not change locations, reflecting their lower willingness to pay for exposed housing. Contrary to the

Table 8: Summary table of relative damage w/ and w/o/ climate change

	Business as usual	Climate change
Inc. group 1		
Inc. equiv. damages	0.67%	1.34%
Exposed pop.	366,265	364,492
Group share	92%	92%
Inc. equiv. loss		*-0.67%*
Inc. group 2		
Inc. equiv. damages	0.12%	0.23%
Exposed pop.	166,655	166,511
Group share	97%	97%
Inc. equiv. loss		*-0.12%*
Inc. group 3		
Inc. equiv. damages	0.10%	0.20%
Exposed pop.	270,802	270,786
Group share	91%	91%
Inc. equiv. loss		*-0.10%*
Inc. group 4		
Inc. equiv. damages	0.07%	0.13%
Exposed pop.	128,747	128,642
Group share	79%	79%
Inc. equiv. loss		*-0.07%*
All		
Inc. equiv. damages	0.32%	0.64%
Exposed pop.	932,469	930,431
Group share	0.91%	0.90%
Inc. equiv. loss		*-0.32%*

Source: Authors' simulations. Each panel gives, separately for each income group, expected flood damages as a share of equivalent net income, the corresponding exposed population and the group share it represents, across simulations without and with added flood risks from climate change. An additional row gives the income equivalent loss obtained from damage increase when introducing climate change.

Figure 11: Relative pluvial flood damage distribution by income group under climate change

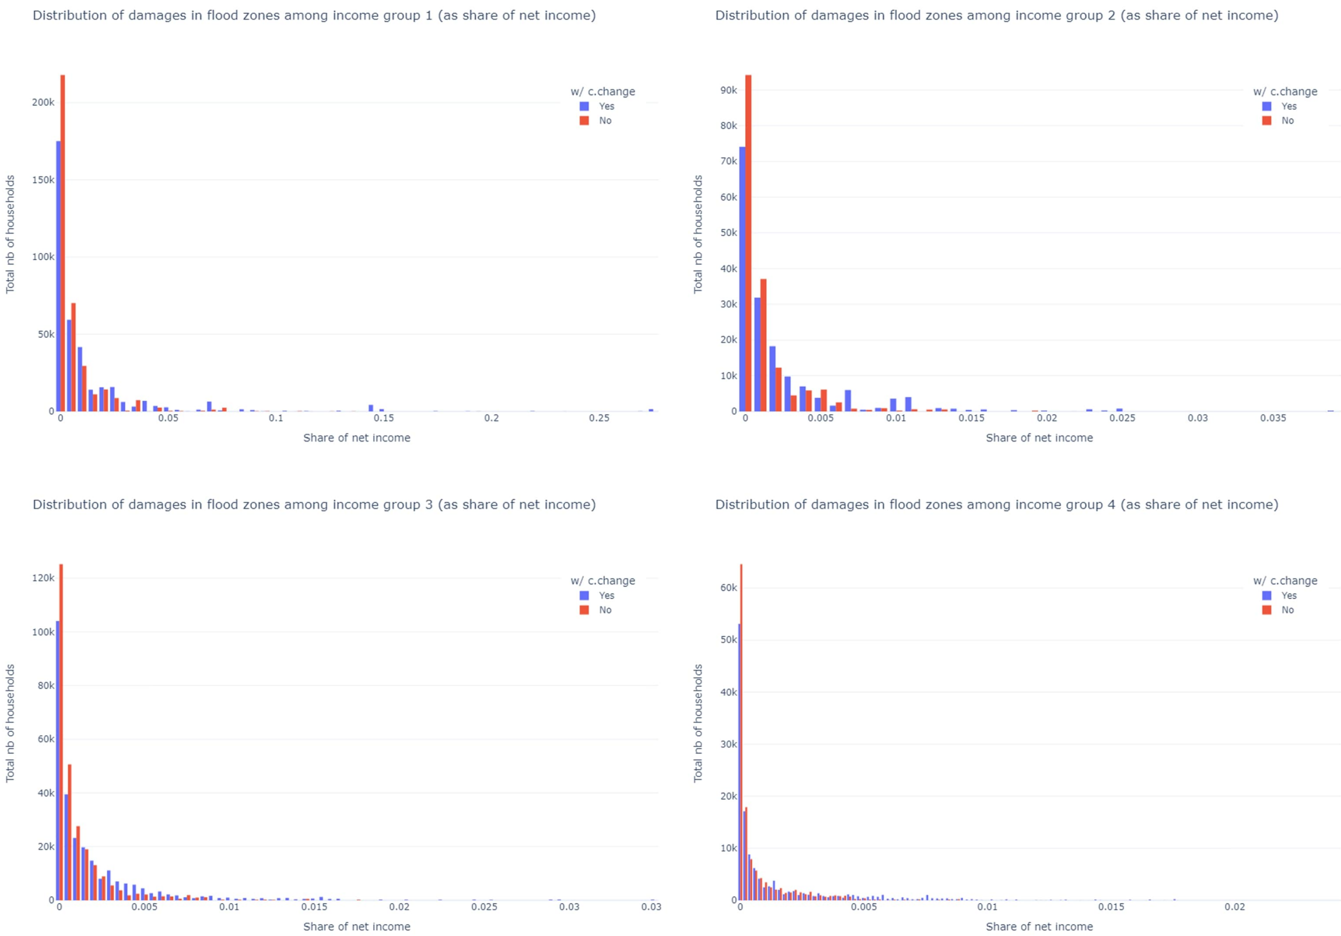

Source: Authors' simulations. Each panel shows the distribution of expected pluvial flood damages across the population as a share of equivalent income net of commuting costs, separately for each income group. Blue bars stand for simulation results with climate change, and red bars for simulation results without climate change.

previous simulation comparison, however, rents mostly fall across the city, with very few rent increases in equilibrium. It is also worth noting that most of the changes are now concentrated in informal housing, which is relatively more vulnerable to pluvial floods due to the absence of drainage systems that we incorporated in risk assessments for the formal housing segment (see Section 3).

All in all, households seem to be responsive to the increase in flood risks due to climate change, but such adaptation at the individual level hides substantial economic damages at the aggregate level.

7 Discussion and extensions [WORK IN PROGRESS]

The previous section made the case for public flood awareness policies for damage reduction. However, to be more realistic, our simulations would have to account for the level of (imperfect) private information at the disposal of poor households especially, and how costly it is for them to adapt their beliefs and behavior.

Figure 12: Spatial population distribution w/ and w/o/ climate change

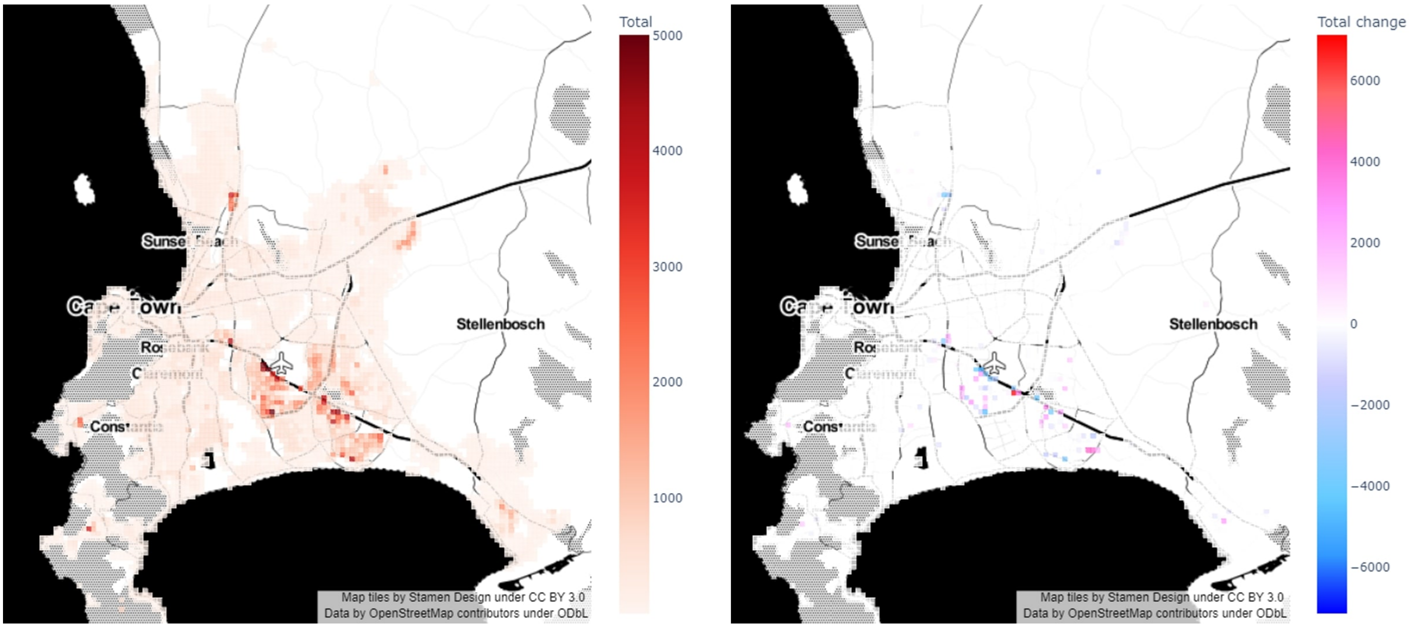

Note: The map on the left shows the values at the benchmark (perfect anticipation). The map on the right shows the change (in absolute terms) when considering climate change.

Besides, it remains to be seen how insurance policies or protective investments may allow them to mitigate the non-negligible share of damages that remains even under perfect information. In doing so, whether private or public efforts should be encouraged depends on the importance of moral hazard, the size of externalities, and the relative cost efficiency of service provision.

To complete this paper, we therefore aim at answering these questions by further augmenting the model and making use of new survey evidence from the City of Cape Town. Our final goal is to provide concrete policy recommendations that may extend to disaster risk management more generally.

8 Conclusion

In this paper, we presented a realistic urban simulation model for low- and middle-income countries, featuring income heterogeneity and informal housing. The key novelty is that we introduced housing damages from three types of flood risks (coastal, pluvial, and fluvial). We estimated the model in the context of the city of Cape Town, South Africa, one of the most exposed cities to pluvial and fluvial flooding in Sub-Saharan Africa. We then simulated the adaptation strategies of households (through relocation and the trade-off between housing and non-housing consumption) following an information shock under current flood risks. We also simulated adaptation strategies following revised flood risks stemming from climate change.

We found that information is effective in helping poor households reduce the adverse consequences of flood risks, notably in the face of climate change notably, in spite of their restricted access to formal housing. In the process, we show that distinguishing between flood types is instrumental in understanding households' responses as households are differently affected by these risks and therefore respond differently. We aim at extending our framework to cover insurance and protective investments, and suggest the most efficient policy approach to disaster risk management.

Figure 13: Absolute change in spatial rent (rands/m²) per housing type in the climate change case (compared with no climate change)

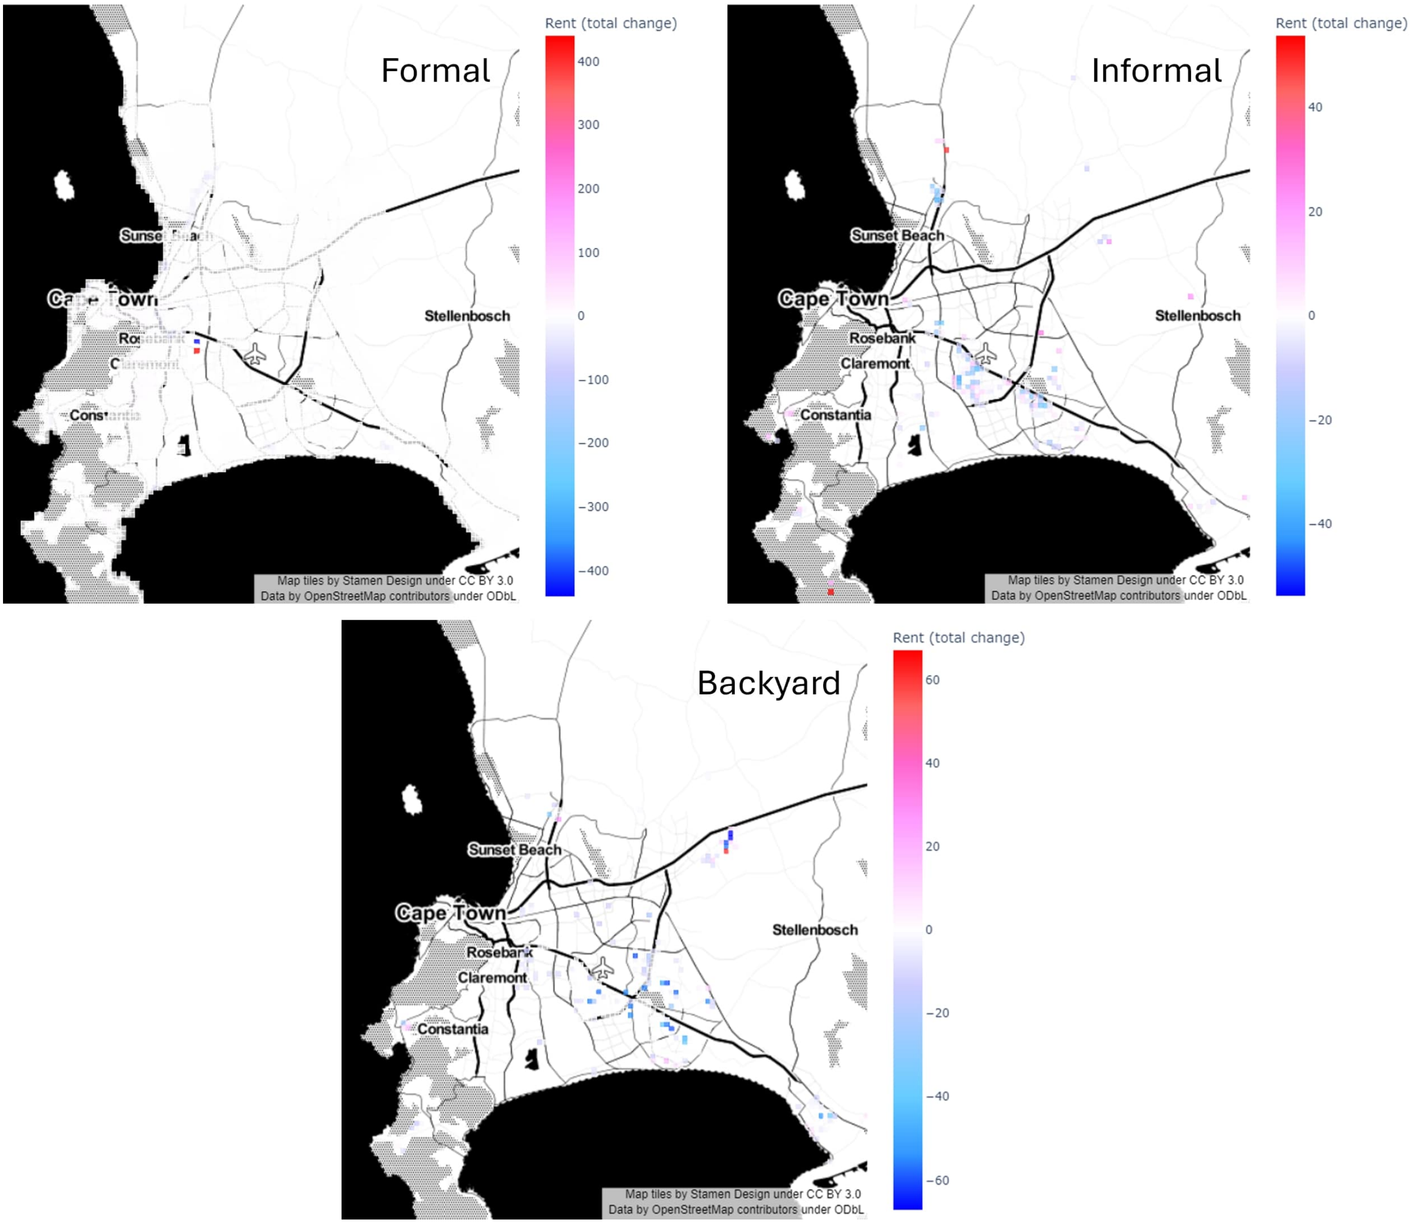

Note: Each map represents by how many rands/m² the market rent would increase (in red) or decrease (in blue) in our simulation with climate change (compared with no climate change), separately for each housing type.

References

M. Allaire. Socio-Economic Impacts of Flooding: A Review of the Empirical Literature. *Water Security*, 3:18–26, 2018.

P. Avner, V. Viguié, B. Arga Jafino, S. Hallegatte, and P. Avner. Flood Protection and Land Value Creation – Not all Resilience Investments Are Created Equal. *Economics of Disasters and Climate Change*, 6(3):417–449, 2022.

C. Balboni. In Harm's Way? Infrastructure Investments and the Persistence of Coastal Cities. *American Economic Review*, 115(1):77–116, 2025.

J. K. Brueckner and H. Selod. A Theory of Urban Squatting and Land-Tenure Formalization in Developing Countries. *American Economic Journal: Economic Policy*, 1(1):28–51, 2009.

J. K. Brueckner, C. Rabe, and H. Selod. Backyarding: Theory and Evidence for South Africa. *Regional Science and Urban Economics*, 79(103486), 2019.

T. Carleton, E. Duflo, K. B. Jack, and G. Zappalà. Adaptation to Climate Change. In *Handbook of the Economics of Climate Change*, volume 1, chapter 4, pages 143–248. 2024.

H. Cole, V. Delbridge, M. d. M. Gómez Ortiz, N. Tsivanidis, and R. D. Zárate. Cities Spatial Model: A Quantitative Spatial Equilibrium Model to Guide Urban Planning, Policy, and Infrastructure Investment Decisions. *International Growth Centre*, 2025.

P.-P. Combes, G. Duranton, and L. Gobillon. The Production Function for Housing: Evidence from France. *Journal of Political Economy*, 129(10):2766–2816, 2021.

G. De Villiers, G. Viljoen, and H. Booysen. Standard Residential Flood Damage Functions for South African Conditions. *South African Journal of Science and Technology*, 26(1):26–36, 2007.

Deltares. Planetary Computer and Deltares Global Data: Flood Hazard Maps. 2021.

I. G. Ellen and R. Meltzer. Heterogeneity in the Recovery of Local Real Estate Markets after Extreme Events: The Case of Hurricane Sandy. *Real Estate Economics*, 52(3):714–752, 2024.

J. Englhardt, H. De Moel, C. K. Huyck, M. C. De Ruiter, J. C. Aerts, and P. J. Ward. Enhancement of Large-Scale Flood Risk Assessments Using Building-Material-Based Vulnerability Curves for an Object-Based Approach in Urban and Rural Areas. *Natural Hazards and Earth System Sciences*, 19(8):1703–1722, 2019.

D. Epple, B. Gordon, and H. Sieg. A New Approach to Estimating the Production Function for Housing. *American Economic Review*, 100(3):905–24, 2010.

J. Finlay and T. C. Williams. Housing Demand, Inequality, and Spatial Sorting. *Journal of International Economics*, 158(104171), 2025.

M. Fujita. *Urban Economic Theory: Land Use and City Size*. Cambridge University Press, 1989.

B. Gonzales, F. Cian, E. S. Nyamador, K. Wienhoefer, and L. Carrera. Living on the Water's Edge: Flood Risk and Resilience of Coastal Cities in Sub-Saharan Africa. *World Bank*, 2023.

S. Hallegatte, C. Green, R. J. Nicholls, and J. Corfee-Morlot. Future Flood Losses in Major Coastal Cities. *Nature Climate Change*, 3:802—-806, 2013.

S. Hallegatte, J. Rentschler, and J. Rozenberg. Lifelines: The Resilient Infrastructure Opportunity. *World Bank*, 2019.

M. J. Hammond, A. S. Chen, S. Djordjević, D. Butler, and O. Mark. Urban Flood Impact Assessment: A State-of-the-Art Eeview. *Urban Water Journal*, 12(1):14–29, 2015.

J. Huizinga, H. De Moel, and W. Szewczyk. Global Flood Depth-Damage Functions: Methodology and the Database with Guidelines. *JRC Technical Report*, 105688, 2017.

A. Kocornik-Mina, T. K. McDermott, G. Michaels, and F. Rauch. Flooded Cities. *American Economic Journal: Applied Economics*, 12(2):35–66, 2020.

Y. Lin, T. K. McDermott, and G. Michaels. Cities and the Sea Level. *Journal of Urban Economics*, 143(103685), 2024.

C. Liotta, V. Viguié, and Q. Lepetit. Testing the Monocentric Standard Urban Model in a Global Sample of Cities. *Regional Science and Urban Economics*, 97(103832), 2022.

C. Liotta, P. Avner, V. Viguié, H. Selod, and S. Hallegatte. Climate Policy and Inequality in Urban Areas: Beyond Incomes. *Urban Climate*, 53(101722), 2024.

B. Marx, T. M. Stoker, and T. Suri. There Is No Free House: Ethnic Patronage in a Kenyan Slum. *American Economic Journal: Applied Economics*, 11(4):36–70, 2019.

B. Merz, H. Kreibich, R. Schwarze, and A. Thieken. Assessment of Economic Flood Damage. *Natural Hazards and Earth System Sciences*, 10(8):1697–1724, 2010.

T. Monnier, B. Pfeiffer, C. Rabe, H. Selod, V. Viguié, and P. Avner. Housing the Urban Poor: Evidence from a Simulation Model with Informality. 2026.

M. Mosimann, L. Frossard, M. Keiler, R. Weingartner, and A. P. Zischg. A Robust and Transferable Model for the Prediction of Flood Losses on Household Contents. *Water*, 10(11): 1596, 2018.

N. Z. Muller and C. A. Hopkins. Hurricane Katrina Floods New Jersey: The Role of Information in the Market Response to Flood Risk. *NBER Working Paper*, 25984, 2019.

F. Ortega and S. Taṣpınar. Rising Sea Levels and Sinking Property Values: Hurricane Sandy and New York's Housing Market. *Journal of Urban Economics*, 106:81–100, 2018.

D. L. Paterson, H. Wright, and P. N. Harris. Health Risks of Flood Disasters. *Clinical Infectious Diseases*, 67(9):1450–1454, 2018.

R. Pharoah. Built-in Risk: Linking Housing Concerns and Flood Risk in Subsidized Housing Settlements in Cape Town, South Africa. *International Journal of Disaster Risk Science*, 5 (4):313–322, 2014.

N. Picarelli, P. Jaupart, and Y. Chen. Cholera in Times of Floods: Weather Shocks and Health in Dar es Salaam. *International Growth Centre*, 2017.

S. Redding and E. Rossi-Hansberg. Quantitative Spatial Economics. *Annual Review of Economics*, 9:21–58, 2017.

J. Rentschler, P. Avner, M. Marconcini, R. Su, E. Strano, M. Vousdoukas, and S. Hallegatte. Global Evidence of Rapid Urban Growth in Flood Zones since 1985. *Nature*, 622(7981):87–92, 2023.

L. Rice, A. Scheba, and A. Harris. Renting in the Informal City: The Role of Dignity in Upgrading Backyard Dwellings in Cape Town, South Africa. *Journal of Modern African Studies*, 61(2):209–233, 2023.

Y. Roux. A Comparative Study of Public Transport Systems in Developing Countries. Master's thesis, University of Cape Town, 2013.

C. C. Sampson, A. M. Smith, P. B. Bates, J. C. Neal, L. Alfieri, and J. E. Freer. A High-Resolution Global Flood Hazard Model. *Water Resources Research*, 51(9):7358–7381, 2015.

D. Sturm, K. Takeda, and A. Venables. How Useful are Quantitative Urban Models for Cities in Developing Countries? 2023.

C. Teulings, I. Ossokina, and H. L. F. de Groot. Land Use, Worker Heterogeneity and Welfare Benefits of Public Goods. *Journal of Urban Economics*, 103:67–82, 2018.

A. Varela Varela. Surge of Inequality: How Different Neighborhoods React to Flooding. 2023.

V. Viguié and S. Hallegatte. Trade-Offs and Synergies in Urban Climate Policies. *Nature Climate Change*, 2(5):334–337, 2012.

V. Viguié, S. Hallegatte, and J. Rozenberg. Downscaling Long Term Socio-Economic Scenarios at City Scale: A Case Study on Paris. *Technological Forecasting and Social Change*, 87: 305–324, 2014.

World Bank. South Africa Country Climate and Development Report. 2022.

World Bank. Strengthening Flood Resilience in Rapidly Growing Cities. 2026.