The Geography of Life: Evidence from Copenhagen*

Gabriel M. Ahlfeldt
Humboldt University Berlin

Ismir Mulalic
Copenhagen Business School

Caterina Soto Vieira
London School of Economics

Daniel M. Sturm
London School of Economics

21 July 2025

Preliminary and Incomplete Draft

Abstract

We use newly constructed individual-level panel data for Denmark covering more than 30 years to document and understand the origins and consequences of spatial sorting by age and family status within cities. Based on by reduced-form evidence that disentangles sorting by age and family status from correlated individual effects, we develop a quantitative urban model with heterogeneity by skill, age, and family status. We use the quantified model to show that the main mechanism driving spatial sorting is heterogeneity in preferences for local amenities. Finally, we use the estimated model to explore the effect of demographic change — population aging, falling fertility, and increasing singlehood — on cities and find that their combined effect on cities will be substantially more muted than any of these changes on their own.

JEL Codes: R21, R23, J12, J13, J14.
Keywords: Spatial sorting, quantitative spatial models, aging, fertility.

*We would like to thank Dave Donaldson, Tatjana Kleineberg, Xavier Jaravel, and seminar participants at University of Barcelona, Collaborative Research Center Rationality and Competition, Copenhagen Business School, Danish Center for Social Science Research (VIVE), German Federal Statistical Office, Harvard/MIT Seminar in International Economics, Helsinki GSE, University of Heidelberg, Humboldt University, IfW Kiel, LSE, University of Luxembourg, Pompeu Fabra, Rockwool Foundation Berlin, and University of Toronto for comments and suggestions. Daniel Sturm gratefully acknowledges financial support from the European Research Council (Grant 101020068 - COMMUTE). We would like to thank many research assistants, especially Carolyn Diamond and Alexander Hansen, for outstanding work. The usual disclaimer applies. Ahlfeldt: Department of Economics, Humboldt University, and Berlin School of Economics, Germany, email: g.ahlfeldt@hu-berlin.de; Mulalic: Department of Economics, Copenhagen Business School, Denmark, email: imu.eco@cbs.dk; Soto-Vieira: Department of Economics, London School of Economics, UK, email: c.soto-vieira@lse.ac.uk; Sturm: Department of Economics, London School of Economics, UK, email: d.m.sturm@lse.ac.uk.

1 Introduction

Many outcomes – such as income, skills, age, fertility, and marriage rates – differ markedly across space, both between cities and within them. However, it remains unclear whether these spatial patterns reflect fixed differences across individuals or cohorts or systematic changes in location choices that relate to life circumstances such as age and family status. We know even less what mechanisms are driving this sorting. Why do young adults, for example, tend to live centrally while families prefer suburban areas? These choices could be due to heterogeneity in preferences for housing consumption, commuting, or local amenities such as entertainment options or access to parks and water areas.

This paper uses newly-constructed employer-employee-property-family panel data for the population of Denmark covering over 30 years to provide evidence on how location choices evolve with age and life events. We document several new stylized facts on how workplace and residence choices, commuting, and residential space consumption change with age and in response to life events. Our data allows us to identify a wide range of life events, including cohabitation, childbirth, separation, empty nesting, retirement, and death of a spouse. We combine the reduced-form evidence with a quantitative urban model to examine what mechanisms, such as differences in housing expenditure shares, commuting costs or the evaluation of residential amenities drive the strikingly different sorting of different groups in space. Finally, we use the estimated model to assess the effect of key demographic changes such as population aging, declining cohabitation rates and fertility on the geography of cities through model counterfactuals.

The Danish registry data are ideal to study spatial sorting at a micro-geographic scale. In our annual panel data for the entire population of Denmark, we observe the residence location and, for those who are working, also the workplace location to a precision of 100 meter by 100 meter grid cells. The data allow us to link people from year to year and observe their marital status, children, income, education, and the size and value of their primary residence. This wealth of information allows us to chart how location choices develop over the life cycle in response to life events and age at a fine geographical scale. This contrasts with information from typical census data where it is difficult to link individuals across censuses and location choices are observed at lower temporal frequency and for much larger geographical units.

We begin by establishing four new stylized facts about how location choices vary over the life cycle. First, we show that both the probability of moving residence and the probability of moving workplace reach a peak of nearly 40% at the age of 21 and then continuously decline with age. Strikingly, most of this mobility is within cities. Over 80% of residence moves and nearly 90% of workplace moves occur within the Copenhagen Metropolitan Area, with similar patterns in other Danish commuting zones. This motivates our focus on mobility across neighborhoods within a

city rather than the much more infrequent mobility across cities. Second, we show that people move closer to the city center until their mid 20s, after which they gradually shift outward, with this decentralization mostly completed by age 40, and much more pronounced for residence than workplace location. Third, commuting times increase by about 10% from the mid 20s to mid-40s with minimal further increases after this age and a small decline just before retirement age. Quantitatively the variation in commuting times is much smaller over the life cycle than variation in the distance of the residence or workplace from the center of the city. Fourth, we show that the average residential floor space consumption per adult continuously increases until the mid 40s and then plateaus for roughly 15 years before increasing further. We therefore find no evidence for "downsizing" later in life.

An obvious question is whether the changing location choices and floor space consumption over the life cycle can be explained by life events such as marriage and childbirth. To do so we estimate event-study specifications using the imputation estimator proposed by Borusyak et al. (2024) and Liu et al. (2024). We estimate the effect of a broad set of life events on distance of residence and workplace from the center of the city, commuting times and the consumption of residential floor space. For the non-working population, we only observe their residential location and residential floor space consumption as outcome variables. We consider all life events that are experienced by at least 5% of the individuals in our sample.[1] This selection rule results in 12 life events, which we group into a set of early and late life events. The early life events are the first three cohabitations, the birth of the first three children, and the first two separations, which all have a median age at which these events occur of less than 35. The late life events are empty nesting, retirement and death of a spouse, which have median ages at event of over 50. We separate our estimation between early and late life events, which assumes orthogonality between these two sets of events.[2] We jointly estimate the event studies, to account for the potential correlation between life events. We do so by using the imputation method to estimate the effects of all early life events jointly in one regression, and we further decompose this total effect into the leads and lags of each event.

The event-study estimates reveal a number of striking patterns. First, cohabitation substantially reduces floor space consumption per capita and decentralizes both the residence and workplace location even though we are controlling for the effects of other life events such as children. Second, the arrival of the first child also decentralizes the residence and workplace location along with a

[1] More than 5% of individual in our sample experience the birth of their third child, for example, while the birth of fourth and further children is experienced by less than 5% of the individuals in our sample and these events are therefore not included in the regressions.

[2] The imputation estimator uses observations before treatment to estimate the person and age fixed effects. Even with 34 years of data it is rarely possible to observe a person that is untreated with all early life events who also experiences a late life event during the sample period. For this reason, we separately estimate the effect of the early and late life events, which implicitly assumes there is no correlation between events that occur in these different stages.

reduction in commuting time. The arrival of the first child also increases residential floor space consumption, but this effect is surprisingly small relative to the space savings from cohabitation. Third, the effects of a separation are almost exactly the opposite of cohabitation. A separation results in both the residence and workplace location becoming more central again, with reductions in commuting time and much larger residential floor space consumption per capita. Fourth, we find hardly any effect of empty nesting or retirement on location choices and floor space consumption. Finally, the death of a partner has minimal effects on the residence location of the surviving spouse but increases space consumption per capita substantially.

A natural question is how much of the changes in residence and workplace location, commuting times, and floor space consumption over the life cycle is explained by life events relative to the effect of age. To decompose the observed changes over the life cycle between age and life events, we use our event-study treatment effects and compute for each individual and year, the joint effect of all life events that they experience. We compute the effect of age as the residual outcome once we remove individual fixed effects and the joint effect of all life events. We find that while life events can explain a substantial part of the variation over the life cycle in residential location choices, they have less of an effect on the workplace location, commuting times and floor space consumption. This suggests that changes in preferences with age that cannot be attributed to observable life events have a considerable impact on location choices within cities.

From the reduced form evidence alone, it is difficult to tell why people of different ages and subject to different life events make such strikingly different location choices within cities. These differences in location choices could be due to several factors. With children, housing needs may grow, making locations with lower house prices relatively more attractive. Similarly, the opportunity cost of commuting time may depend on the circumstances of life, and those with lower commuting costs may sort into peripheral locations with longer commutes. Finally, the valuation of location-specific amenities such as nearby bars, restaurants or schools may change with age and family structure. To disentangle the importance of these factors we develop a quantitative urban model with agents who are heterogeneous in terms of their age, marital status, skill level and the presence of children. We take advantage of the comprehensive micro data available in Denmark to estimate the structural parameters of this model for a uniquely rich set of worker and family types. In particular, we consider a model with 12 different types of workers.

We use the estimated model to evaluate the relative importance of different channels for the strikingly different location choices of different groups. We do this through a series of model counterfactuals where we eliminate different mechanisms through which the model can explain the observed sorting of groups. Before we present the full model counterfactuals, we examine the direct impact effect of key parameters on sorting, which are first-order approximations to the full general equilibrium effect of changes in these parameters. We find that groups that on average

live further from the center of the city actually have lower rather than higher housing expenditure shares, which suggests that high housing expenditure shares do not explain sorting into cheaper suburban locations. Similarly, we show that groups that on average prefer more suburban residential locations in fact have higher commuting costs than those living more centrally, which suggests that lower commuting costs induce groups to choose more peripheral locations. Finally, we show that differences in amenities across locations are highly predictive of the average distance of a group's residence location from the center of the city, which points to the central importance of amenities in understanding location choices within cities.

The full counterfactual results confirm the pattern of results from the first-order approximations. The first model counterfactual assumes that all agents have the same housing expenditure share and finds that the average distance of a group's residential location from the center of the city barely changes. Further counterfactuals look at the importance of commuting costs, idiosyncratic utility shocks, labor demand and residential amenities. We find that eliminating the variation in each of these parameters across groups has very limited effects on the average distance of a group's residential location from the center of the city apart from the effect of variation in amenities. If all groups perceived the amenity value of a location to be the same, sorting of groups across locations would dramatically change.[3] These results suggest that differences in the perceived residential amenity value of a location for different groups are key for understanding location choices.

We next correlate the model estimated amenities by group with observables proxies for amenities that have been widely used in the literature. In particular, we create a consumption amenity index, which is composed of the density of cinemas, restaurants, bars and street food establishments; a natural amenity index, which includes a dummy for location on the coast, the share of water, green areas, the share of non-built-up land and the number of marinas. We further construct a school quality index, based on test scores and completion rates. We correlate these indices with the model amenities by group, and interact them with being in our older age group, cohabiting, having children or being high-skilled. The results show that, relative to our reference group (young, low-skilled single people without children), workers who are older, in a couple, and cohabiting with children value consumption amenities less, and natural amenities more. Parents and senior workers value school quality relatively more than the reference group. These results suggest that the model inverted amenities of each group are intuitively correlated with variation in observable urban amenities across locations and reinforces our finding that amenity differences are key to understanding sorting patterns in cities.

In a final step, we use the quantitative model to assess the impact of key demographic changes

[3]Formally, we assume in the amenity counterfactual that each location has the same residential amenity value for all groups, and we set this common amenity value equal to the unweighted average amenity level of all groups with non-zero population in a location in the baseline.

on the geography of cities. In particular, we estimate model counterfactuals for population aging, changes in fertility (i.e. the number of people with children) and the share of single households. We find that each of these demographic trends has substantial impacts on location choices and house prices within cities. However, their combined effect is more muted. While population aging reduces demand for central areas, this is offset by the effect of fewer children and more single households, which increases demand for space in the center of cities.

Related Literature Our paper is related to a number of literatures. First, our paper is related to the literature that has examined the effect of life events, such as marriage, childbirth, divorce, and retirement, on economic outcomes other than location choices. One branch of this literature investigates the effect of the first child on earnings and labor force participation rates (such as Kleven et al. (2019), Adda et al. (2017), and Cortés and Pan (2023)), and consumption choices (such as Browning and Ejrnæs (2009)). Kleven et al. (2024) jointly estimates the effects of marriage and children on labor market participation rates across the world. Foerster (2024) examines the impact of divorce on hours worked, wages, assets, labor income, and consumption, and Fernández and Wong (2014) shows that increased probability of divorce impacts female labor force participation. We contribute to this literature by studying how life events affect location choices, which are of first-order importance due to the large share of housing expenditure for most individuals and large differences in house prices, amenities, and labor markets across locations. Moreover, our data allows us to jointly estimate the impact of different life events and evaluate their relative importance.

Second, our work is related to the large literature that studies how economic outcomes evolve with age. This includes work explaining labor supply over the life cycle starting with seminal contributions by Ben-Porath (1967), Mincer (1974), and Heckman (1976). Another prominent strand in this literature investigates how consumption and savings vary over the life cycle following the seminal contribution by Modigliani and Brumberg (1954). Recent surveys on this literature include Meghir and Pistaferri (2010). This literature has typically ignored location choices and housing consumption despite the large share of housing in typical consumption baskets and the large variation in amenities and the price of housing across space. We contribute to this literature by showing how age affects location choices and housing consumption.

Third, our paper builds on the literature developing quantitative urban models that capture the determinants of location choices with cities, including Lucas and Rossi-Hansberg (2002), Ahlfeldt et al. (2015), Allen et al. (2015), Heblich et al. (2020), Owens III et al. (2020), Gaubert and Robert-Nicoud (2024), Couture et al. (2024), Tsivanidis (2023), Redding and Sturm (2024), Bordeu (2024) and Weiwu (2024), as recently surveyed in Redding (2024).[4] We build on the strand

[4]There is also a large body of literature on horizontal residential sorting that focuses solely on the housing market

of this literature that considers how heterogeneous agents sort in cities. Due to our rich data, we can estimate the structural parameters of a model with many worker types, far beyond the small number typically considered in the existing literature, which usually distinguishes workers based only on skills. Furthermore, we contribute to this literature by investigating how the structural residuals of our model – the residential amenities – are correlated with real-world amenities, and how these differ by group characteristics.[5]

Fourth, there is a recent literature that examines how life events such as retirement and child birth or age affect locations choices. Komissarova (2022) and Badilla Maroto et al. (2024) look at migration across cities in response to retirement in the US and France respectively.[6] Moreno-Maldonado and Santamaria (2024) show how delayed childbirth affects sorting across neighborhoods in US cities while Anstreicher and Venator (2024) find that young mothers are more likely to move back to their US state of birth. Coeurdacier et al. (2024) use French data to show that the number of children is higher in locations with lower house prices both within and across cities and rationalize this finding in an overlapping generations model. Albouy and Faberman (2025) estimate amenity differences across US metropolitan areas and find that particularly the high-skilled move to higher amenity cities early on in their life using data from the National Longitudinal Surveys of Youth. We contribute to this emerging literature in several ways.

First, our panel data allows us to control for individual fixed effects to disentangle cohort effects from the treatment effect of life events and age on location choices. Second, we jointly estimate the effects of a broad set of life events and age and assess their relative importance. Third, our quantitative model allows us to explore the importance of different mechanisms such as commuting costs, housing expenditure shares and amenities in explaining the patterns of spatial sorting.

The remainder of this paper is organized as follows. Section 2 provides an overview of our data. Section 3 presents our stylized facts for location choices in cities. Section 4 introduces our quantitative model. Section 5 describes the estimation of the model and how we use it to shed light on the mechanisms for sorting. Section 6 uses the estimated model to explore the effects of demographic changes on cities, and Section 7 concludes.

---

to explain household location choices, linking household taste heterogeneity to local urban amenities. For an overview of residential sorting models, see Kuminoff et al. (2013). The methodology used in these models was originally developed by Berry (1994), Bayer et al. (2007) and Bayer and Timmins (2007).

[5] Relatedly, Ang et al. (2024) investigates how much real world amenities explain the variation of model amenities in Los Angeles, California, derived from a canonical urban quantitative spatial model with one group.

[6] A further related paper is Bonnet et al. (2010), which shows that widowhood increases the probability of changing residence location using French data.

2 Data

We use a newly compiled individual-level panel dataset covering the entire population of Denmark for the years 1986 to 2019. This employer-employee-property-family dataset contains information on residence and workplace locations down to 100 by 100 meter grid cells, income from all sources, residential floor space consumption and property prices as well as detailed demographic information including age, education, marital status, and children. The dataset includes both the working population as well as the non-working population such as retirees, students, and children. We provide an overview of the data in this section with more detail on the construction of each variable presented in Online Appendix Section A.

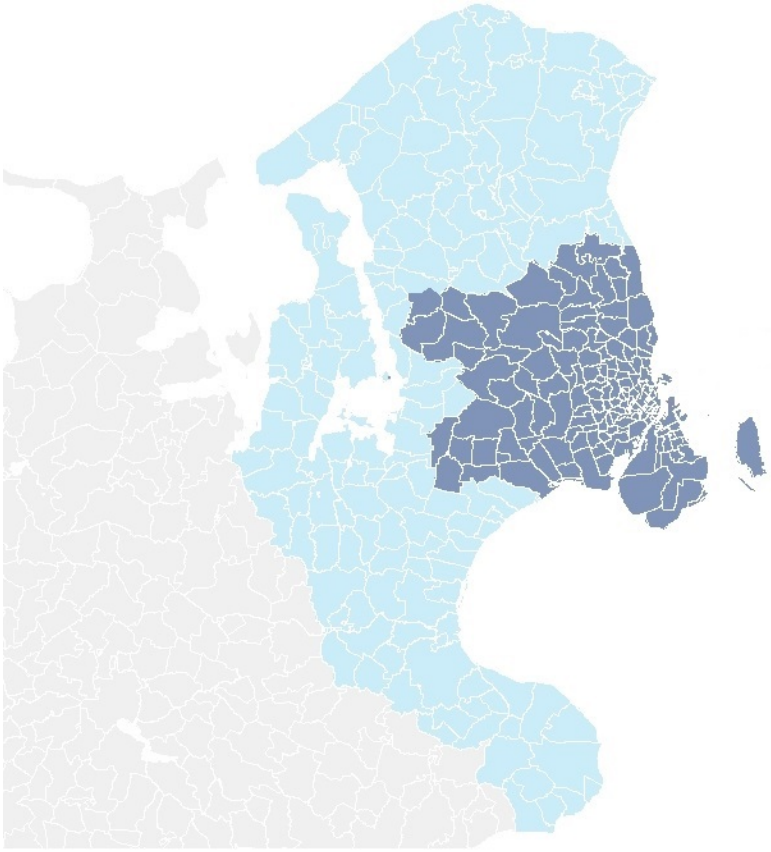
While our dataset covers the entire population of Denmark, most of our results use those living and working in Copenhagen, which is the capital and largest city of Denmark. Figure 1 shows a map of Copenhagen. Our primary definition of the city is the Copenhagen Metropolitan Area (CMA), which is the combined light blue and dark blue areas of Figure 1 and covers 3,038 square kilometers, including even distant suburbs of Copenhagen and substantial areas of open space that are not built up.[7] In 2015, the total residential population of the CMA was just under 2 million inhabitants and 92 percent of those working in this area also lived in it. In robustness checks, we also examine the Greater Copenhagen Area (GCA), shown in dark blue in Figure 1. Using this narrower definition of the city, we find that 79 percent of people working in the area in 2015 also lived within this area.

Residence and Workplace Location and Income At the beginning of each year, Statistics Denmark records the residential address for each individual and also the workplace address for each employed individual residing in Denmark.[8] The recorded workplace corresponds to the individual's primary work location rather than the firm's head office in cases where firms operate across multiple locations. To protect the confidentiality of individual level data, Statistics Denmark does not disclose the precise coordinates of residence and workplace addresses, but we are able to observe them in 100 by 100 meter (i.e. one hectare) grid cells. For each person, we observe their entire work and non-work income in each year and their hours of work as well as their occupation. We also observe the year individuals retire and begin receiving pension benefits, as well as the amount of these pension payments. For retired individuals or those not working for other reasons (such as full-time students), we only observe their place of residence.

[7] The north-south and east-west extension of the Copenhagen Metropolitan Area are 99 km and 63 km, respectively.

[8] Housing policy in Denmark strongly discourages individuals from occupying multiple residential units, with only a limited number of holiday homes exempt from this policy.

Figure 1: Copenhagen Metropolitan Area (CMA)



Notes: The figure shows the spatial extent of the Copenhagen Metropolitan area (light blue and dark blue areas combined) and the Greater Copenhagen area (dark blue area). The borders shown in white indicate parish boundaries.

Residential Floor Space Consumption We observe the universe of residential units in Denmark in each year and are able to link each residential unit to those living in the unit. The data on residential units includes information on the size of the unit in square meters, its type (such as whether it is a flat or house), and the age of the building. We use this data to compute residential floor space consumption per person by dividing the total floor area of a residential unit by the number of adult residents living there. The number of adults excludes children below the age of 18 and also children aged 18 to 25 who are still living with their parents.[9] We observe whether people own or rent their residence. In 2011, just under 45% of those living in the CMA were owner-occupiers, about 10% live in social housing and the remainder rent on the private market.

[9] A relatively small number of adult children in Denmark live with their parents. Not counting adult children as adults in the household for the purposes of floor space consumption per person avoids mechanical jumps in household floor space consumption per person when children turn 18 or move away from home shortly after turning 18.

Property Prices For the years 1992 to 2019, we observe the universe of real estate transactions in Denmark, including each transaction price and exact location.[10] We also observe a number of property characteristics, such as the size of the unit in square meters, the type of unit (multi-family or single-family house), and the age of the building. We use the algorithm developed by Ahlfeldt et al. (2023) to estimate house price indices for our spatial units. For the years 1986 to 2011, we also have access to tax assessments of the value of each property in Copenhagen, which were prepared by local property assessors and reported to Statistics Denmark.[11]

Family Structure and Children For each person, we observe in each year whether they are single or cohabiting with a partner. A large majority of those cohabiting are either married or living in a registered civil partnership. For those not legally married or in a civil partnership, Statistics Denmark uses a set of rules to determine whether people living at the same address are a couple or are just flat-sharing.[12] We also observe when one spouse in a couple passes away. Finally, for each individual, we observe whether they have children, the birth dates of these children, and whether the children still reside in the parental household.

Education For each person, we observe their completed level of education. We aggregate educational attainment to either high or low skill by assuming that those with no more than vocational education are low-skilled while those with at least a bachelor degree are high-skilled. We also observe when people are currently in full-time education.

Travel Times To reduce the computational burden of calculating travel times, we aggregate the 100 by 100 meter grid cells into larger 200 by 200 meter cells. There are 30,150 such 200 by 200 meter cells in the Copenhagen Metropolitan Area that have at least one resident or one person working in this cell at some point during our sample period. For each of these approximately 909 million (30,150 $\times$ 30,150) bilateral connections, we calculate travel times in 2015 in minutes. We calculate travel times by car, bicycle, and walking using an ArcGIS shape file of the street network. For public transport, we determine minimum travel times using the metro, suburban railway, and bus networks. To compute the fastest connection by public transport, we allow passengers to combine all modes of public transport and walking to minimize travel time. We measure overall travel times by weighting public transport, walking, cycling, and car minimum travel times, using

[10] Real estate transaction data for Denmark prior to 1992 is unavailable, as digital registration only began in that year. Earlier transactions were recorded manually in physical archives and have not been digitized.

[11] All property owners in Denmark pay a property tax based on the assessed value of their property. Property assessments have, however, been suspended since 2011, while a new property assessment system is being developed. The 2011 assessments have been used as the tax base for the property tax since 2011.

[12] While Denmark has allowed same-sex marriages since 2012, the numbers are small and we do not distinguish between same-sex and heterosexual couples.

travel zone-level data on the proportion of journeys undertaken with each mode from the Danish National Travel Model described in Rich and Overgaard Hansen (2016).

3 Stylized Facts

In this section we provide evidence on how individual's location choices evolve over the life cycle and to what extent these choices are due to life events. In Section 3.1, we show how moving probabilities evolve with age and the relative importance of within city versus across city mobility. In Section 3.2 we show how location choices, commuting time and residential floor space consumption vary over the life cycle. In Section 3.3, we quantify the contribution of key life events, such as cohabitation and having children, to the life cycle in location choices. Finally, in Section 3.4 we estimate the relative contribution of life events and age to the overall life cycle pattern in location choices.

3.1 Mobility over the Life Cycle

We begin by using our data to document how often people move either their place of residence or their workplace. As discussed in Section 2, we observe peoples' residence and workplace to a precision of 100 by 100 meter grid cells at an annual frequency and count as a move a change in location across these grid cells between one year and the next. We restrict attention to those over the age of 18 up to 65 for workplace moves, and up to the age of 90 for residence moves, as the data are sparse above these ages.[13] Figure 2a shows the frequency of residential and workplace moves for those living in Copenhagen as a function of age. The figure shows a striking pattern with residential and workplace mobility peaking at the early age of 21. After this early peak, mobility rapidly declines with age, with workplace mobility decreasing slower than residential mobility and remaining higher than residential mobility.

Figure 2b explores to what extent people move between different locations within Copenhagen or from Copenhagen to other parts of Denmark or abroad. The figure shows that the vast majority of moves occur within cities, both in terms of residence and in terms of workplace changes. In particular, the figure shows that over 80% of residential moves and nearly 90% of workplaces moves are across locations within Copenhagen. Moreover, this pattern remains relatively stable

[13] As discussed in more detail in Online Appendix Section A, we use the following conventions in counting moves. We do not count the first job of a person as a workplace move. As we restrict attention to those 18 and older, the first time someone can change their place of residence in our data is therefore when they are 19. We observe both inflows into and outflows from unemployment and only count the outflow as a workplace move. A workplace move from abroad to Denmark is treated like a first job and does not count as a workplace move. By construction, this measure of mobility ignores if someone makes multiple residence or workplace moves in a single year or those who move such short distances that both their origin and destination are in the same 100×100 meter grid cell.

Figure 2: Residence and Workplace Mobility over the Life Cycle

(a) Probability of moving residence or work



(b) Conditional probability of moving within CMA



Notes: This figure presents two panels describing patterns of individual mobility over the life cycle by age. Panel (a) plots the probability of moving either residence or workplace as a function of age. Panel (b) displays the conditional probability that a residence or workplace move occurs within the CMA, plotted against age.

over the life cycle. There is a small drop in within Copenhagen residential mobility right after the age of 60, indicating that some people may move out of Copenhagen after retiring. After this short decline, the probability of staying in Copenhagen increases steadily with age.

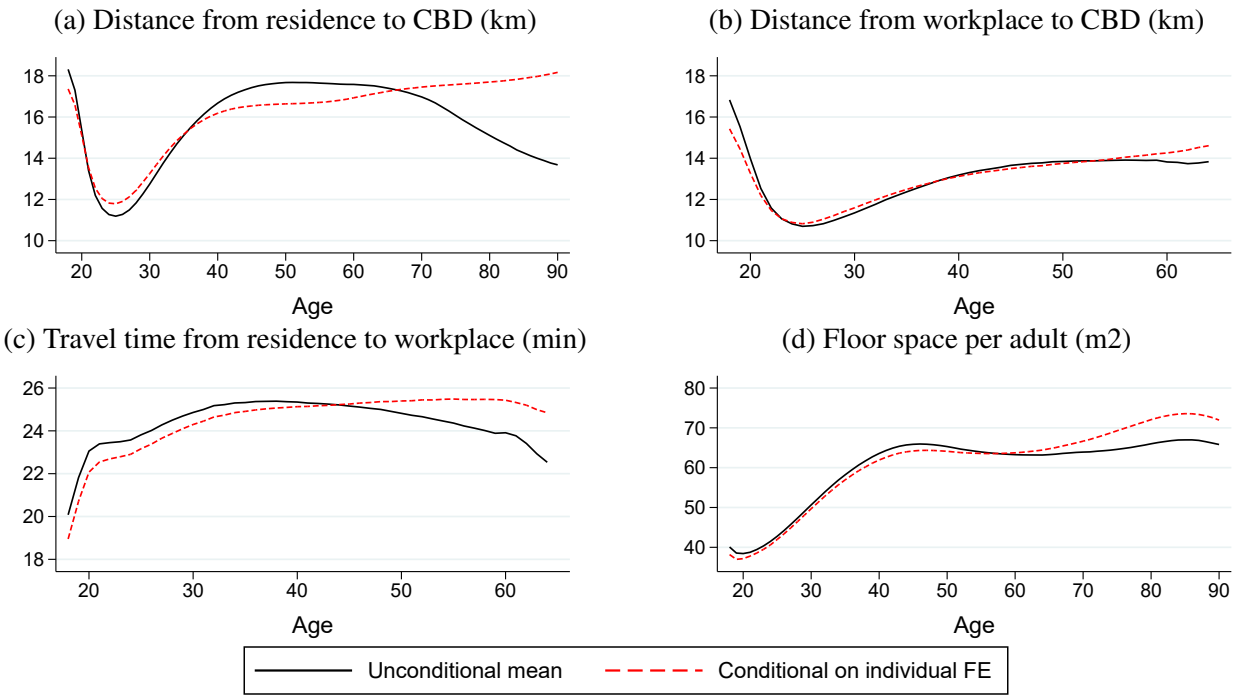
3.2 Location Choices over the Life Cycle

In this section, we show how key outcomes related to the geographical choice of individuals change over their life cycle. Using the panel of individuals, we can follow residence and workplace location choices across an individual's age. For each individual, we compute the average distance between the residence and center of the city and, if they work, the average distance between their workplace and the center of the city. We observe the size of the residential units where individuals live, and the number of adults with which they share a household. We therefore compute a measure of floor space per adult in square meters.

In Figure 3a, we display the average distance in kilometers from the residence location to the center of the city by age. The black solid like shows the unconditional mean, while the red dashed line shows the conditional mean, after controlling for individual fixed effects. The figure shows that early in the adult life, people on average live further away from the center, but soon centralize their residences. The inflection point is at age 26, when a process of decentralization starts. The differences between the two lines shows that accounting for individual fixed effects is important, especially at older ages, due to differences in survival probabilities particularly at higher ages.

In Figure 3b, for the individuals who work, we compute the mean distance in kilometers of their workplace from the center of the city and plot it against age. Due to most individuals being retired after age 65, we cap the line at that age. Individuals start their career further away from the

Figure 3: Residence and Workplace Location Choices over the Life Cycle



Notes: This figure shows how residence and workplace location choices evolve with age across four dimensions. Panel (a) shows the distance from residence to the CBD. Panel (b) shows the distance from workplace to the CBD. Panel (c) reports travel time from residence to workplace in minutes, while Panel (d) shows floor space per adult in square meters. Each panel plots the unconditional mean in black and values conditional on individual fixed effects in dashed red.

center of the city, and soon start centralizing their jobs. There is also a pattern of decentralization of the workplace after the mid-20s, but this pattern is much less stark than that in residential location. Furthermore, there are small differences between the unconditional mean and the mean conditional on individual fixed effects, with the gap only increasing when people approach their late 50s. This is likely showing that individuals that retire earlier work in difference places from those that choose to retire later.

Using the bilateral travel times described in Section 2, we measure the average travel time in minutes from the residence location to the workplace for each individual in our sample and plot it against age in Figure 3c. The pattern in this panel mirrors the patterns of the previous two panels of this figure. People start their work life commuting shorter distances since both residence and workplace are on average decentralized. We see a relatively small increase in the average commuting time compared to how much individuals centralize both residence and workplace. The commuting time in minutes increases when individuals move further out to the suburbs, and this decentralization is only partially accompanied by a decentralization of workplace.

Lastly, we compute the average floor space per adult in each household and observe how this evolves over the life cycle in Figure 3d.[14] We observe that individuals start their adult life consuming small amounts of floor space in square meters, but this amount increases consistently until individuals are about forty years old. This increase in floor space consumption could be due to the formation of families and parenting. In later stages of the life cycle, the mean conditional on individual fixed effects increases further. There is no evidence of "downsizing" as people reach higher ages, apart from a slight drop in space consumption after age 85, which could be due to moves to care homes after this age. In contrast, we observe that the unconditional mean remains relatively constant suggesting that the consumption of floor space per adult for individuals that survive in our sample is downward biased, relative to the sample average.

3.3 Life Events

The results of Section 3.2 suggest that location choices and floor space consumption have clear life cycle patterns. Life events, such as marriage and having children, could play an important role in driving these observed trends. In this section, we explore if important life events, such as getting married and having children, can at least partially explain the changes in location choices and floor space consumption over the life cycle.

We observe when individuals cohabit, when they separate, when they have children, when children leave home, and when they become widowed, among other important life milestones. For consistency, we define life events by the order in which they occur (first separation, second cohabitation, etc).[15] In our data, we find substantial variation in the age at which different life events are experienced, and if people experience them at all (see Online Appendix Table C1). There is also considerable variation in the order in which individuals experience such events, although naturally some events imply a specific order (one cannot have a second child without having had the first already).[16] Moreover, we consider that some of these events are correlated; individuals that have a first child may already consider a second one, couples that start living together may already consider having a child, etc. However, events that occur in later stages of life are less likely to be correlated with early life events: becoming a pensioner is plausibly not influenced by when and if you have had a first or second separation. We split our analysis into early and late life events. Furthermore, we only consider life events that happen to at least 5% of the individuals in our sample. The full list of life events is in Online Appendix Table C1.

We estimate the effects of the multiple life events jointly using the methodology of Borusyak

[14] Following the definition of Statistics Denmark, we do not consider as adults children between 18 to 25, who are still living with their parents.

[15] See Online Appendix Section A for more details on the construction and definition of the variables used to determine the life events.

[16] Online Appendix Figures C9 and C10 show the frequency of the events by age.

et al. (2024), and extend it for estimating multiple events jointly in the same regression. The details of the econometric model are presented in Online Appendix Section B. The estimation of the event studies is based on a counterfactual estimator, where the estimation of treatment effects is separated from the estimation of the age and individual fixed effects (Borusyak et al. 2024, Liu et al. 2024). This method prevents the issue of the forbidden comparisons and negative weights that arise when using two-way fixed effects estimators (Callaway and Sant'Anna 2021, Sun and Abraham 2021). Moreover, this method is more efficient than using two-way fixed effects estimators, greatly reducing estimation time.

The estimation procedure for our counterfactual estimator is as follows. First, we estimate the age and individual treatment effects with the untreated sample (untreated and yet-to-be treated). We consider untreated observations as those who have not yet been treated by *any* life event of this life stage.[17] The key identifying assumption is that, in the absence of treatment, the trajectory of outcomes would have remained unchanged. We proceed by imputing the estimated age and individual fixed effects on the outcome variable for all of the sample, thus predicting the treatment effect of life events for each observation. We jointly estimate the average treatment effect of each life event by decomposing the imputed outcome on the leads and lags of each life event in a single step. This ensures that the correlation across life events does not generate omitted variable bias. We describe our method and lay out the necessary assumptions in Online Appendix Section B.

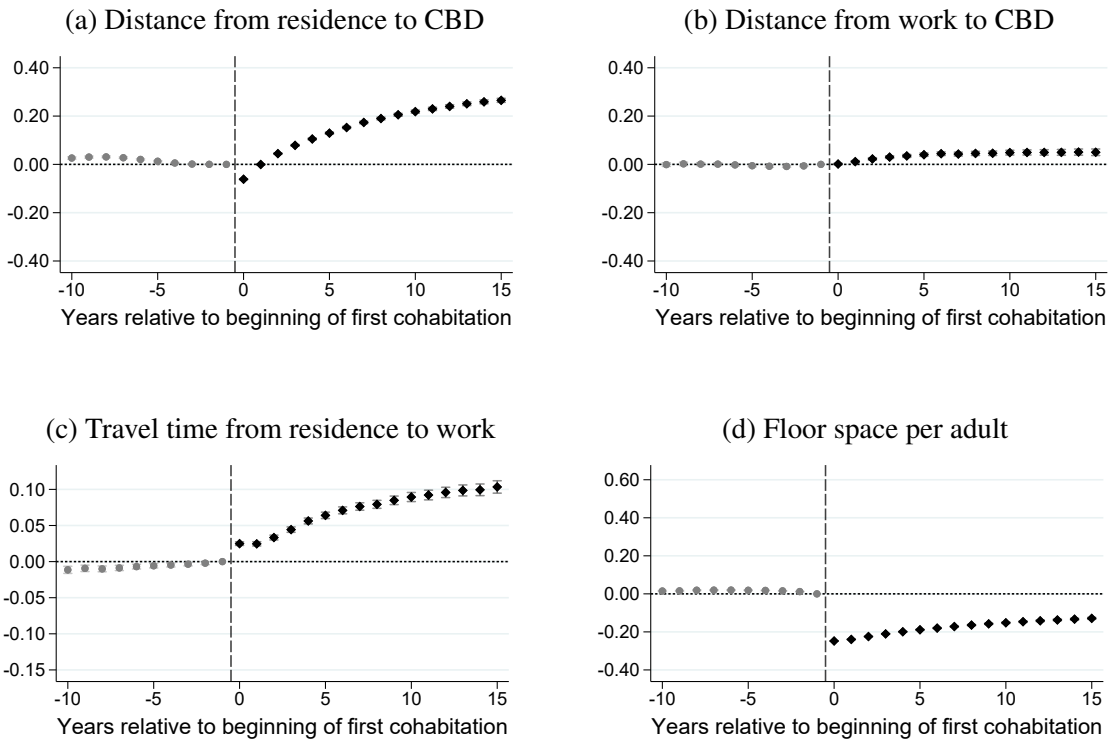
Now we show a simplified and concise version of our method in one equation only. Consider a set of life events $e \in \mathbb{E}$, for which we want to estimate the average treatment effects. Equation 1 shows the event-study regression in a single step:

$$Y_{it} = \hat{\theta}_i + \hat{\eta}_a + \sum_{e \in \mathbb{E}} \sum_{\substack{h=-m \\ h \neq -1}}^{n} \beta_h^e \mathbb{1}[K_{it}^e = h] + \varepsilon_{it} \tag{1}$$

where Y_{it} is the outcome of individual i in year t. The individual and age fixed effects are represented by $\hat{\theta}_i$ and $\hat{\eta}_a$, respectively, estimated previously with the untreated sample. $K_{it}^e = t - E_i^e$ is the difference between the current year (t) and the year in which individual i experiences event e (E_i^e), and $\mathbb{1}[K_{it}^e = h]$ is a dummy for difference h. The treatment effects of the n leads and m lags of the life event $e \in \mathbb{E}$ are denoted by β_h^e. The regressions contain all leads and lags but we restrict the graphs to show the treatment effects from -10 to $+15$. We consider two different sets of events $\mathbb{E}$ for our regressions, early and late life events. Implicitly, we assume that these sets of events are uncorrelated, which, considering the distance in ages between them seems a reasonable assumption.

[17] We consider two life stages: early and late life. Therefore, for the early life events one needs to not have been treated with cohabitation or children, while for the late life events, the untreated sample needs to not have been treated with retirement, empty nesting and widowhood.

Figure 4: Effects of First Cohabitation



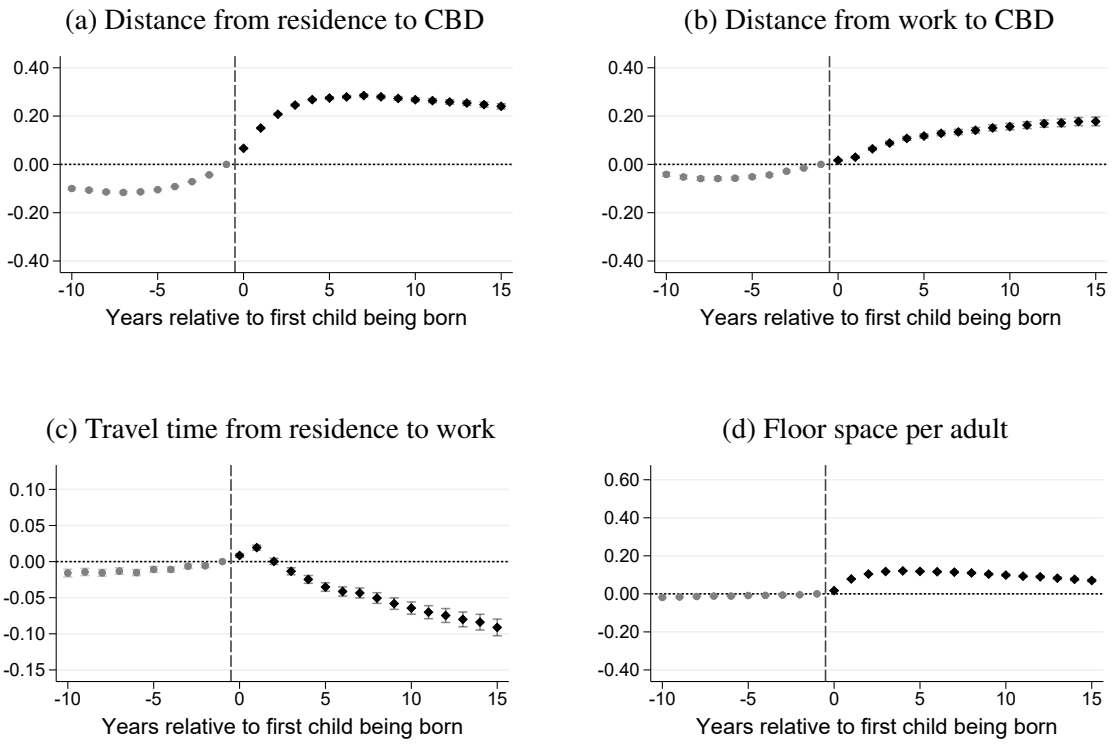
Notes: Each panel of this figure plots event-time coefficients relative to the beginning of the first cohabitation at year 0. Panel (a) shows the change in distance from residence to the CBD, and Panel (b) shows the change in distance from workplace to CBD in kilometers. Panel (c) reports the changes in commute time from residence to workplace in minutes. Panel (d) shows the change in floor space per adult in square meters.

We look at the effects of life events on four main outcomes: distance from residence to the center of the city, distance from work to the center of the city, commuting time from home to work, and floor space per adult in square meters. Across all graphs showing the effects of any event on the same outcome, we keep the scale of the y-axis fixed, so that we can easily compare the relative importance of each life event.

Figure 4 shows the effect of the first cohabitation, which is the first life event for most of our sample.[18] The figure plots the point estimates and the confidence intervals, but due to the large sample size, the standard errors are barely visible. Figure 4a shows that the first cohabitation has a large and persistent effect on decentralizing individuals which increases over time, despite a small centralization effect in the first year of cohabitation. Figure 4b shows that there is also a significant decentralizing effect on the workplace location, but it is much smaller in comparison to the effects shown in Figure 4a. Consequently, we observe a relative increase in travel time

[18]The sample used to estimate the early life events contains 1,323,393 unique individuals. Approximately 87% of the individuals cohabit before their first child, while 11.4% experience both in the same year. Figure C8 shows how the majority of our sample experiences first cohabitation before having a first child.

Figure 5: Effects of First Child



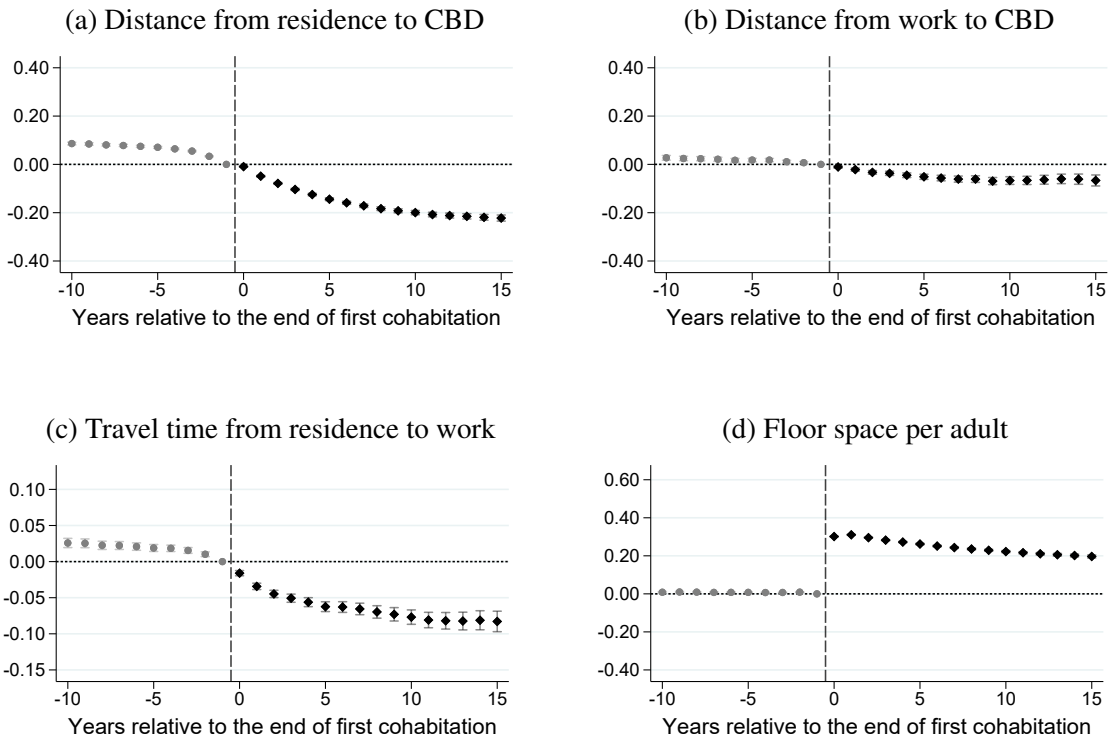
Notes: The four panels of this figure plot event-time coefficients relative to the birth of the first child at year 0. Panel (a) shows the change in distance from residence to the CBD, and Panel (b) shows the change in distance from workplace to the CBD in kilometers. Panel (c) reports the change in commuting time from residence to work in minutes. Panel (d) shows the change in floor space per adult in square meters.

between residence and workplace in Figure 4c. Notably, Figure 4d shows that when individuals start cohabiting, they decrease their consumption of floor space per adult by about 25%, which illustrates that cohabitation rates are an important driver of residential floor space demand in cities.

Figure 5 analyzes the effect of having the first child on location decisions. We estimate these event studies relative to the year before the mother was likely pregnant and not relative to the year before we observe the child as a member of the household.[19] The two top panels suggest that there are some anticipation effects with respect to both residence and workplace location of having a child. Figure 5a shows that households suburbanize after the firstborn arrives, and this effect is large and persists for many years after the child is born. Figure 5b shows that individuals also move their workplace location in response to the first child, with their employment location becoming less central. Figure 5c shows that commuting time decreases on average after the first child is born, with a small increase in the first two years followed by a considerable decrease in commuting time

[19]Most data in our panel is from January of each year. If we observe a new child in January of a year, the mother must have been pregnant for at least part of the previous year and may already have been pregnant in year $t-2$. We define the year before a child appears as $t=0$, but some mothers will already have been pregnant in $t-1$.

Figure 6: Effects of First Separation



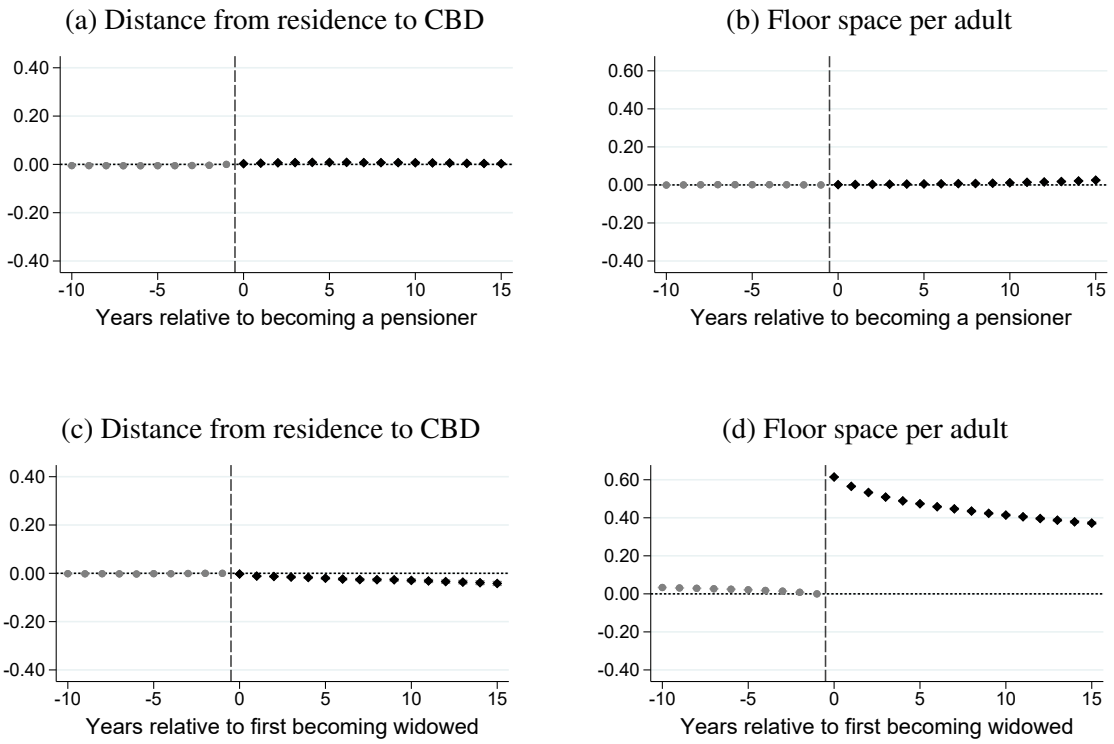
Notes: The four panels of this figure plot event-time coefficients relative to the first separation at year 0. Panel (a) shows the change in distance from residence to the CBD, and Panel (b) shows the change in distance from workplace to the CBD in kilometers. Panel (c) reports the change in commuting time from residence to work in minutes. Panel (d) shows the change in floor space per adult in square meters.

the following years. Figure 5d shows how expanding the family increases the demand for floor space, with floor space per adult increasing by approximately 10% after individuals have their first child. Despite there being anticipation effects of children on both residence and workplace location, we do not see any anticipation effects in floor space consumption.

In Figure 6, we show the estimates of the effects of the first separation on the location choices and floor space consumption of individuals. We consider separation as the end of a cohabitation, when individuals stop living with each other. Figure 6a indicates that individuals slightly anticipate separation, and start moving more centrally even before they decide to stop living together. Figure 6b shows that the individuals respond by changing workplace right after getting separated, choosing on average more central employment locations. Figure 6c shows a decline in commuting time after a separation, a result of centralizing both residence and workplace. Figure 6d shows that after separation, individuals increase their consumption of floor space per adult by approximately 35%.

Across all panels, we observe that the first separation nearly perfectly mirrors the effects of the first cohabitation, with opposite results in terms of direction. These findings can be due to different

Figure 7: Effects of Retirement and Widowhood



Notes: The four panels of this figure plot event-time coefficients relative to retirement and widowhood at year 0. Panels (a) and (c) show the change in distance from residence to the CBD, and Panels (b) and (d) show the change in floor space per adult in square meters.

preferences, with single people valuing amenities present in the center of the city, such as bars and restaurants, and potentially a more efficient marriage market.[20] Furthermore, while not exhibited here, the estimates for second separation are strikingly similar to those present in Figure 6, despite much fewer people ever experiencing a second separation. In Table C1, we observe that while 27.2% of the sample experiences a first separation, only 8.2% experience a second.

We now present the results with the treatment effects of the late life events, which are estimated separately from the early life events.[21] We only estimate these effects on the distance from residence to the CBD and on floor space per adult; outcomes unrelated to work. We do that because, when individuals retire, they mechanically don't have a workplace anymore. Since we jointly estimate the effects of all three late life events, we also do not obtain the effects of empty nest and widowhood on workplace related outcomes.

The top two panels of Figure 7 show the effects of becoming a pensioner on the distance from residence to the CBD and on floor space consumption. We reiterate that we keep the scale

[20] Gautier et al. (2010) show that urban areas tend to attract more singles because they offer more efficient marriage markets due to density.

[21] See Online Appendix B for the details of our methodology.

of the y-axis constant across all figures, so that we can easily compare the relative importance of life events on changing individuals' outcomes. In Figures 7a and 7b, we estimate very small treatment effects of becoming a pensioner, especially when we compare them to the effect size of first cohabitation and first childbearing. Similarly, Figure 7c shows that we estimate relatively small treatment effects of widowhood on distance from residence to the CBD, with individuals that become widowers slightly centralizing their residence location over the course of 15 years. On the other hand, Figure 7d shows that widowhood increases floor space per adult by approximately 60%. A large part of this effect is mechanical, as by construction floor space per adult would increase if there is one fewer adult in the household.

3.4 Decomposing the Effects of Age and Life Events

We have now shown that location choices have a clear age profile, and that life events can affect individuals' location choices. We now proceed to decompose by how much the changes in location choices that occur over life are due to the effect of life events and how much is due to aging alone. The red dashed line in Figure 3 shows the age profile of the outcomes of interest, conditional on individual fixed effects and averaged over age bins. Let $\bar{\mathcal{Y}}_s$ represent the life cycle effects for outcome y_{it}:

$$\bar{\mathcal{Y}}_s = \mathbb{E}\left(y_{it} - \alpha_i \mid \text{Age}_{it} = s\right) \tag{2}$$

We will decompose $\bar{\mathcal{Y}}_s$ into two effects: the joint treatment effect of all the life events we have estimated, $\bar{\mathcal{L}}_s$, and the effect of aging, $\bar{\mathcal{A}}_s$, which are obtained as a residual. Using the coefficients estimated in Section 3.3 to obtain, for each individual i at each point in time t, a value for the total treatment effect of all life events:[22]

$$\hat{\mathcal{L}}_{it} = \sum_{e} \sum_{h=-a}^{b} \widehat{\beta}_h^e \mathbb{1}[K_{it}^e = h] \tag{3}$$

which can be averaged across age bins to obtain:

$$\bar{\mathcal{L}}_s = \mathbb{E}(\hat{\mathcal{L}}_{it} \mid \text{Age}_{it} = s) \tag{4}$$

The aging effects, $\mathcal{A}_{it}$, can be obtained by substracting the effect of life events for person i in time t from the outcome variable net of individual fixed effects at the individual-year level. We can further average the aging effects by year, and we obtain the following relationship:

$$\bar{\mathcal{A}}_s = \bar{\mathcal{Y}}_s - \bar{\mathcal{L}}_s \tag{5}$$

[22] Despite the early and late life events regressions being estimated separately, we use the coefficients from both estimations in this decomposition exercise. This way we capture the total treatment effect of life events.

These curves are normalized to zero at the age of 18, when individuals enter the sample. Therefore, we look at the variation in levels from the baseline level of the outcome when people are 18 years old.

Figure 8 shows the outcomes of our decomposition exercise. The red dashed line is replicated from Figure 3, and represents the effects of the life cycle, $\bar{\mathcal{Y}}_s$, conditional on individual fixed effects. This line is decomposed into a dark blue solid line, representing the effect of life events, $\bar{\mathcal{L}}_s$, and a light gray solid line, representing the effects of aging, $\bar{\mathcal{A}}_s$. Figure 8a decomposes the distance from residence to the CBD, and it shows that the pattern of centralization is counteracted by the force of life events, pulling individuals to locations further away from the city center. It is also clear that during the first years of adulthood, life events have close to no effect in explaining residence location choices of individuals. We can also observe that after the age of forty, the contribution of life events is stable, in accordance with late life events having a much lower contribution to residence location decision.

In Figure 8b, we observe two patterns: i) life events have a relatively smaller contribution on where you work than aging, and ii) in the absence of life events, individuals barely decentralize their workplace. The results of Figure 8c summarize the previous two results. We find that the life cycle in the travel time from residence to work is mostly composed of aging effects, but it is amplified by the effect of life events, which further increase commuting time for individuals. Lastly, we look at what happens to floor space per capita over the life cycle in Figure 8d. While the increases in floor space per capita occur mostly due to aging effects, we can observe a contribution of life events, especially after the age of 30. This contribution decreases slightly, but life events seem to matter throughout the life cycle, even if less than aging.

We have now discussed a series of stylized facts that help us understand how location choices and the consumption of floor space change over the life cycle, and what is the relative contribution of important life events in shaping these outcomes. We now will use a quantitative spatial model to understand how determinants of location choice change across different life cycle groups.

4 Theoretical Framework

In order to understand what drives the spatial sorting of individuals of different ages and life stages, we develop a quantitative spatial model in which groups differ in terms of their age and family type. The model builds closely on the literature following Ahlfeldt et al. (2015), but due to the richness of our data, we can model a much larger number of different groups than the existing literature, which has typically considered two or four types of agents.

We consider a city with a discrete number of locations, $n \in \mathbb{N}$. Locations differ in terms of their productivity and amenities and are connected by a transport network that determines travel times

Figure 8: Decomposing of the Effects of Age and Life Events



Notes: This figure displays the decomposition of the life cycle effect into the effects of age and of life events, plotted against age. Panel (a) shows the change in distance from residence to the CBD in kilometers. Panel (b) shows the change in distance from workplace to the CBD in kilometers. Panel (c) reports the change in commute time from residence to workplace in minutes. Panel (d) shows the change in floor space per adult in square meters.

between locations. Firms hire different types of workers and produce a homogeneous tradable good, which we choose as the numeraire. We consider a closed city with fixed numbers of each type of agent. Furthermore, for simplicity we assume that the supply of both residential and commercial floor space in each location of the city is fixed.[23]

In Section 3, we show that life events and aging can have considerable impact on sorting of residence and workplace location, and on floor space consumption. In our model, these differences are captured by the heterogeneity of the working groups, which may be in different life stages, captured by age and family type. Workers can be young or senior professionals, and can be single, cohabiting or cohabiting with children. We also allow workers to differ between high and low skill to ensure that the results of the model are not due to income sorting. This results in a total of

[23] We observe both the price of residential and commercial floor space in each location, which differ considerably within locations due to planning restrictions that limit transitions between residential and commercial use.

12 groups. The aim of the model is to understand what are the economic factors that explain the sorting of individuals in different life stages, through the lenses of a canonical quantitative urban model.

4.1 Preferences

We consider workers who differ in their occupation o and their family type f. An occupation is the combination of age and skill, and we consider four different occupation groups: young and old high-skilled workers and young and old low-skilled workers. We consider twelve different types of workers of, which differ by age, family type, and skill level, i.e. high-skilled couple with children, for example. The indirect utility for worker ω, with occupation o and family type f residing in location n and working in location i depends on their occupation-specific wage w_i^o, the price of residential floor space Q_n, commuting costs $\kappa_{ni}^{of} = t_{ni}^{\phi^{of}}$, group specific location amenities B_{ni}^{of}, and an idiosyncratic amenity draw $z_{ni}^{of}(\omega)$ for each worker:

$$U_{ni}^{of}(\omega) = \frac{B_{ni}^{of} w_i^o z_{ni}^{of}(\omega)}{\kappa_{ni}^{of} (P_n)^{\alpha_{of}} (Q_n)^{1-\alpha^{of}}} \quad 0 < \alpha^{of} < 1. \tag{6}$$

The group-specific amenity B_{ni}^{of} captures the amenities of a worker of occupation o and family type f living in n and working in i. This overall amenity consists of both a residence component ($\mathcal{B}_n^{of}$) and a workplace component ($\mathcal{B}_i^{of}$) so that $B_{ni}^{of} = \mathcal{B}_n^{of} \mathcal{B}_i^{of}$. The idiosyncratic amenities $z_{ni}^{of}(\omega)$ of each worker are independent draws from a Fréchet distribution for each residence-workplace pair and worker group, following $G(z) = e^{-(z)^{-\varepsilon^{of}}}$. The Fréchet shape parameter $\varepsilon^{of} > 1$ controls the dispersion of idiosyncratic amenity draws and we allow the shape parameter to vary across worker groups.[24]

Workers choose their residence-workplace pair to maximize their utility. The properties of the Fréchet distribution allow us to write the probability that a worker from group of chooses to live in n and work in i as:

$$\lambda_{ni}^{of} = \frac{L_{ni}^{of}}{L^{of}} = \frac{\left(B_{ni}^{of} w_i^o\right)^{\varepsilon^{of}} \left(\kappa_{ni}^{of} (Q_n)^{1-\alpha^{of}}\right)^{-\varepsilon^{of}}}{\sum_{k\in\mathbb{N}} \sum_{\ell\in\mathbb{N}} \left(B_{k\ell}^{of} w_\ell^o\right)^{\varepsilon^{of}} \left(\kappa_{k\ell}^{of} (Q_k)^{1-\alpha^{of}}\right)^{-\varepsilon^{of}}}, \quad n, i \in \mathbb{N}, \tag{7}$$

where L_{ni}^{of} denotes the measure of individuals from group of that have chosen this residence-workplace pair; and L^{of} is the total measure of individuals from group of in the city, which we assume to be exogenously determined. The probability that any individual from group of resides

[24] We normalize the Fréchet scale parameter to one since it is isomorphic to the common amenities B_{ni}^{of}.

in location n is determined by:

$$\lambda_{Rn}^{of} = \frac{L_{Rn}^{of}}{L^{of}} = \frac{\sum_{\ell \in \mathbb{N}} \left(B_{n\ell}^{of} w_{\ell}^{o} \right)^{\varepsilon^{of}} \left(\kappa_{n\ell}^{of} \left(Q_n \right)^{1-\alpha^{of}} \right)^{-\varepsilon^{of}}}{\sum_{k \in \mathbb{N}} \sum_{\ell \in \mathbb{N}} \left(B_{k\ell}^{of} w_{\ell}^{o} \right)^{\varepsilon^{of}} \left(\kappa_{k\ell}^{of} \left(Q_k \right)^{1-\alpha^{of}} \right)^{-\varepsilon^{of}}}, \quad n \in \mathbb{N}, \tag{8}$$

4.2 Production

The production of the tradable good occurs under conditions of perfect competition and constant returns to scale. Firms only observe workers' life stage (young or old) and skill (high or low), and not their family circumstances. An equivalent assumption is that firms are not allowed to offer different wages to workers based on their family circumstances. As a result, there are four different wages which are offered to young low-skilled workers, old low-skilled workers, young high-skilled workers and old high-skilled workers. Output is produced with a Cobb-Douglas technology that uses all four types of labor inputs (L_{Fi}^{o}) and commercial floor space (H_{Fi}) as inputs:

$$Y_i = A_i \prod_{o \in \mathbb{O}} \left(\frac{L_{Fi}^{o}}{\beta_i^{o}} \right)^{\beta_i^{o}} \left(\frac{H_{Fi}}{\beta^{H}} \right)^{\beta^{H}}, \quad 0 < \beta_i^{o}, \beta^{H} < 1, \quad \sum_{o \in \mathbb{O}} \beta_i^{o} + \beta^{H} = 1, \tag{9}$$

where A_i is the productivity of location i. The Cobb-Douglas labor input shares (β_i^{o}) are allowed to vary across locations. The following zero-profit condition must hold in each location with positive production of the tradable final good:

$$A_i \prod_{o \in \mathbb{O}} \left(\frac{1}{w_i^{o}} \right)^{\beta_i^{o}} \left(\frac{1}{q_i} \right)^{\beta^{H}} = 1, \quad 0 < \beta_i^{o}, \beta^{H} < 1, \quad \sum_{o \in \mathbb{G}} \beta_i^{o} + \beta^{H} = 1, \tag{10}$$

where q_i is the price of commercial floor space and β^{H} is the input share of commercial floor space.

4.3 Commuter Market Clearing

Commuter market clearing requires that the measure of individuals with occupation o and family type f working in workplace i must equal the measure of individuals from this group choosing to commute to location i from any residence location n:

$$L_{Fi}^{of} = \sum_{n \in \mathbb{N}} \lambda_{ni|n}^{of} L_{Rn}^{of} \tag{11}$$

where $\lambda_{ni|n}^{of}$ is the probability that workers in occupation o and of family type f commute to workplace i conditional on living in residence n:

$$\lambda_{ni|n}^{of} = \frac{\left(\mathcal{B}_i^{of} w_i^{o} / \kappa_{ni}^{of} \right)^{\varepsilon^{of}}}{\sum_{\ell \in \mathbb{N}} \left(\mathcal{B}_{\ell}^{of} w_{\ell}^{o} / \kappa_{n\ell}^{of} \right)^{\varepsilon^{of}}}. \tag{12}$$

Commuter market clearing also implies that the per capita income in residence location n for the worker with occupation o and family type f is a weighted average of the wages in all locations, where the weights are the conditional commuting probabilities by residence $\lambda_{ni|n}^{of}$:

$$v_n^{of} = \sum_{i \in \mathbb{N}} \lambda_{ni|n}^{of} w_i^o. \tag{13}$$

4.4 Floor Space Market Clearing

As we do not model the production process for floor space and assume that both residential and commercial floor space in each location are exogenous, floor space market clearing in each location requires:

$$H_{Ri} = \sum_{o \in \mathbb{O}} \sum_{f \in \mathbb{F}} (1 - \alpha^{of}) \frac{v_i^{of} L_{Ri}^{of}}{Q_i} \tag{14}$$

$$H_{Fi} = \beta_H \left(\frac{A_i}{q_i} \right)^{\frac{1}{1-\beta_H}} \prod_{o \in \mathbb{O}} \left(\frac{L_{Fi}^o}{\beta_i^o} \right)^{\frac{\beta_i^o}{1-\beta_H}} \tag{15}$$

4.5 General Equilibrium

Given the model's structural parameters $\{\phi^{of}, \alpha^{of}, \beta^H, \beta^o, \varepsilon^{of}\}$, the city's populations by group $\{L^{of}\}$, and the exogenous vector of location specific characteristics $\{A_i, B_{ni}^{of}, H_{Fi}, H_{Ri}\}$, the general equilibrium of the model is referenced by the vector of endogenous variables $\{L_{Ri}^{of}, w_i^o, L_{Fi}^{of}, Q_i, q_i\}$ determined by the following equations: Worker residential choice probabilities (8); profit maximization and zero profit condition (10); commuter market clearing for each worker type (11); residential floor space clearing (14); commercial floor space clearing (15).

5 Quantifying the Mechanisms

In this section, we describe how we quantify the structural parameters of our model and use the quantified model to examine the mechanisms behind the sorting of different groups in Copenhagen. For the estimation of the model, we aggregate our data to the level of parishes, of which there are 288 in the Copenhagen Metropolitan Area. This level of aggregation still captures the rich heterogeneity in sorting of different groups across the city, and avoids the issue of too many zeros in the counts of employment and residents in each spatial unit and commuting flows across locations. Online Appendix Section E provides more details on the estimation for each parameter.

The estimation proceeds in the following steps. First, we estimate housing expenditure shares for each group using detailed data on income, house prices and space consumption by each household. Second, we estimate gravity equations for each group in the model using commuter flows between parishes and our travel time matrix. Third, we estimate the Fréchet shape parameter and decompose the gravity coefficient estimates into the Fréchet shape parameter of each group and the elasticity of commuting cost with respect to travel time of each group. Fourth, we calibrate the parameters of the production function across locations using data on the labor force composition in each location. Fifth, we invert the model to recover the residential amenities of each group in each location and the productivity of a location.

In Section 5.1, we discuss the estimation of the model parameters and the model inversion. In Section 5.2, we discuss the relative explanatory power of key model parameters for sorting in the city. Finally, in Section 5.3, we use model counterfactual to determine the full general equilibrium contribution of each model parameter for the observed sorting in the city.

5.1 Parameter Estimation and Model Inversion

In this section, we discuss the estimation of parameters and the model inversion. Online Appendix Section E provides more details. Table 1 summarizes the estimates of the utility function parameters. We start by describing the housing expenditure share, which we calibrate from our data. Column (1) shows our estimates of the housing expenditure share of each of the 12 groups in the model. We estimate the housing expenditure shares using our data on individual income, residential floor space per capita and estimated square meter prices. In particular, we calculate the estimated value of an individuals residential floor space consumption and convert this to an annual fair market rent using an assumed rental return. This estimated rental return is divided by the individual's income to arrive at the housing expenditure share of this individual. We calibrate the rental return of residential units so that the average housing expenditure share in each year is 30%, which is a value estimated by the Danish council of economic advisors. The housing expenditure shares that we estimate with this procedure varies from around 24% for several groups to a peak of about 35% for young singles, which imply considerable differences in exposure to house price differences across groups.

To estimate how sensitive different groups are to travel time, we take the natural logarithm of (7) after substituting $\kappa_{ni}^{of} = t_{ni}^{\phi^{of}}$ and obtain the conventional commuting gravity equations of log commuter flows on travel times in minutes:

$$\log \lambda_{ni}^{of} = \gamma_n^{of} + \gamma_i^{of} - \nu^{of} \times t_{ni}^{of} + u_{ni}^{of} \tag{16}$$

where $\nu^{of} = \phi^{of}\varepsilon^{of}$. In this regression, the origin fixed effects (γ_n^{of}) capture the effects of residential amenities and residential floor space prices. The destination fixed effects (γ_i^{of}) account for

the workplace amenities and wages. Columns (2) and (3) of Table 1 show the results of estimating these gravity equation using PPML and PPML IV respectively.[25] The PPML IV estimates use straight-line distance between parishes as an instrument for travel times, which avoids the potential endogeneity in the placement of infrastructure where exogenous travel demand shifters induce planners to invest in the speed of particular connections (Redding 2024). The PPML IV estimates are higher than the PPML estimates, which is consistent with transport planners in Copenhagen steering investments to areas with lower transport demand to promote local economic development. A complementary explanation for the upward bias in OLS is classical measurement error in travel times.[26]

Comparing the estimates for high-skilled and low-skilled workers reveals that low-skilled workers have somewhat larger gravity coefficients in absolute magnitude, which implies that their commuting flows are more sensitive to travel time than those of high-skilled workers. This finding is in line with estimates in Tsivanidis (2023) for Colombia and Heblich et al. (2020) for London. Comparing the estimates for young and old workers reveals no major differences between these two groups. We confirm this impression by estimating a pooled gravity equation for all young and all old workers and cannot reject the null hypothesis that these two gravity coefficients are identical. Furthermore, there are not major differences in the gravity coefficients across family types of singles, cohabiting individuals, and individuals cohabiting with children. Most of the differences that exist across age and family groups when we estimate the model with PPML disappear when we use the IV, suggesting that they were, at least in part, driven by measurement error or omitted variable bias.

With the gravity coefficients at hand, we start the model inversion procedure, in which we estimate additional parameters and recover location fundamentals. The model inversion closely follows Ahlfeldt et al. (2015), with adaptations to account for fewer labor market dimensions (o) than worker groups (of). Our departure point is Equation 11, which we use to estimate the Fréchet shape parameter (ε^{of}) of each group of workers. Let the transformed wages for each group be defined as $\omega_i^{of} = (\mathcal{B}_i^{of} w_i^o)^{\varepsilon^{of}}$. With the commuting decays ν^{of} estimated in our gravity equations for each worker group, we can rewrite the commuter market clearing condition as follows, where the only unknown variable is the transformed wages, which are identified up to scale:

$$L_{Fi}^{of} = \sum_{n \in \mathbb{N}} \frac{\omega_i^{of} / t_{ni}^{\nu^{of}}}{\sum_{\ell \in \mathbb{N}} \omega_\ell^{of} / t_{n\ell}^{\nu^{of}}} L_{Rn}^{of}. \tag{17}$$

Using employment and resident population data, we use an algorithm to find the transformed wages that satisfy Equation 17. Let us define the adjusted wages as $\tilde{w}_i^{of} = \mathcal{B}_i^{of} w_i^o = (\omega_i^{of})^{1/\varepsilon^{of}}$.[27]

[25] Even at the level of parish to parish flows, the share of zero flows in all bilateral connections ranges from 74% to just under 87%, illustrating the importance of controlling for zeros in these gravity regressions.

[26] We use PPML IV following Wooldridge (2015), using a control function approach.

[27] Throughout the paper, we use the terms adjusted wages and model implied wages interchangeably.

Table 1: Estimated Model Parameters by Group

Model Group	Housing Expenditure Share	ν (PPML)	ν (PPML IV)	ε	ϕ
Young, single, low-skill	34.7%	−0.080	−0.126	6.009	−0.021
Young, single, high-skill	35.7%	−0.056	−0.097	6.046	−0.016
Young, cohabiting, low-skill	24.4%	−0.064	−0.104	5.924	−0.018
Young, cohabiting, high-skill	25.0%	−0.043	−0.079	7.886	−0.010
Young, cohabiting with children, low-skill	24.2%	−0.077	0.122	9.854	−0.012
Young, cohabiting with children, high-skill	23.7%	−0.060	−0.103	8.045	−0.013
Senior, single, low-skill	34.7%	−0.081	−0.129	6.674	−0.019
Senior, single, high-skill	35.4%	−0.064	−0.106	6.221	−0.017
Senior, cohabiting, low-skill	24.1%	−0.079	−0.126	7.021	−0.018
Senior, cohabiting, high-skill	23.3%	−0.066	−0.109	5.822	−0.019
Senior, cohabiting with children, low-skill	24.1%	−0.076	−0.121	7.683	−0.016
Senior, cohabiting with children, high-skill	24.4%	−0.062	−0.101	5.873	−0.017

Notes: The table shows the estimated parameters by group. We calibrate housing expenditure shares using a computed residential house price index for locations in the CMA, and household consumption of floor space. We compute the gravity coefficients using PPML, and we use straight line distances as an IV for bilateral travel times.

For each group, we estimate ε^{of} through an algorithm that, given transformed wages, searches for the ε^{of} that matches the moment condition that the variance of the log adjusted wages must equal the variance of the log wages in the data. Following this procedure, we obtain values of the Fréchet shape parameter for all the worker groups.

Once both ε^{of} and ν^{of} are estimated, we can obtain the commuting decay elasticity by group (ϕ^{of}). Our paper is the first, to the best of our knowledge, to estimate Fréchet shape parameters for so many different groups, but our estimates are within the range of the estimates from other papers that have used this estimation strategy, such as Ahlfeldt et al. (2015) and Heblich et al. (2020). Our average estimate across the 12 worker groups is 6.56, very close to the estimate of 6.83 of Ahlfeldt et al. (2015). We don't observe systematic differences in ε^{of} across the skill dimension, but we observe considerable variation in the Fréchet shape parameter across age and family type dimensions. On average, young workers display larger estimates than older workers, and families with children have large estimates of ε^{of}. This parameter indicates the importance of the idiosyncratic component of utility, with lower values indicating higher variance. This suggests that the decision of groups with higher values of ε^{of} are more subject to prices and amenities relative to idiosyncratic preferences in comparison to groups with lower estimates. Lastly, we obtain the estimates of ϕ^{of}. We observe that singles, seniors, and low-skilled workers have, on average, larger commuting decay elasticities in absolute values.

We continue to describe the inversion of the model, after estimating ε^{of} and the transformed wages for each group. With these estimates, we can obtain the adjusted wages by group ($\tilde{w}_i^{of}$), identified up to scale. We proceed by inverting the occupation wage (w_i^o) and the workplace amenities for each group ($\mathcal{B}_i^{of}$). In this step, we make the identifying assumption that for each

location the average workplace amenities for groups within an occupation is one. That is, we have that $\mathbb{E}_f(\mathcal{B}_i^{of}) = 1$.

We adjust the occupation wages so that the aggregate wage bill share by occupation in the model matches the data. This assures that differences in wage bill across occupations that exist in the city are maintained in our model. This step allows us to maintain both the differences in aggregate wage levels across occupations and the differences in wages across space within occupations.

In our model, the production function has labor input shares that are occupation and location specific. Therefore, using the model adjusted occupation wages $\tilde{w}_i^{of}$, we compute the occupation wage bill divided by the total wage bill in each location. We rescale these wage bill shares, so they reflect the share of total labor inputs in the production function $(1-\beta^H)$:

$$\beta_i^o = \frac{\sum_f w_i^o L_i^{of}}{\sum_o \sum_f w_i^o L_i^{of}} (1-\beta^H) \tag{18}$$

The floor space input share in the production function (β^H) is calibrated to 15%, using estimates from the Danish Sages.[28] With the production function parameters and the occupation wages, we use the zero-profit condition from (10) to recover location productivities A_i. We estimate the group-specific residence amenities for workers using the residence location choice probabilities detailed in (8). Given the adjusted wages, residence population shares and floor space prices, we are able to back out amenities for all groups. We finish inverting the model by computing the residential floor space quantity that clears the demand from residents, and the commercial floor space quantities that clears the demand from firms.

5.2 First-Order Effects on Sorting

In our data, we observe rich cross-sectional sorting patterns that are representative of many cities around the world: young singles and childless couples are overrepresented in central areas, whereas families with children and older age groups prefer the suburbs.

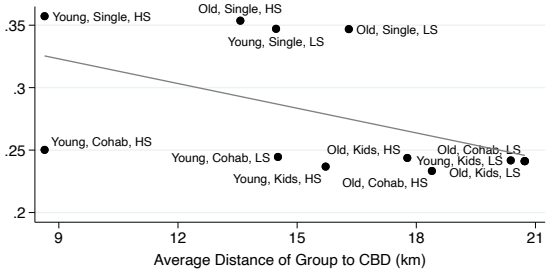
Our reduced-form analysis in Section 3 confirms that this cross-sectional sorting pattern is not purely driven by correlated cohort and individual effects, but is to a large extent the result of preferences that vary depending on family status and a worker's stage of life. Our quantitative model allows us to gain further insights into what type of preferences drive the sorting. Before we quantify the relative contributions of different mechanisms in counterfactuals, we descriptively examine how groups living in different parts of the city differ in terms of the model primitives that determine spatial sorting in Figure 9.

Since housing cost enters indirect utility negatively and is weighted by the expenditure share on housing, groups with a greater housing expenditure share will, all else equal, prefer suburban

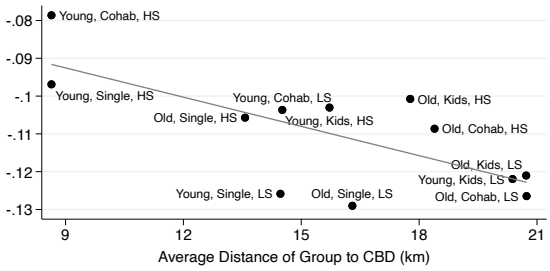
[28] See the 2021 Productivity report in https://dors.dk/vismandsrapporter/produktivitet-2021.

Figure 9: Determinants of Spatial Sorting

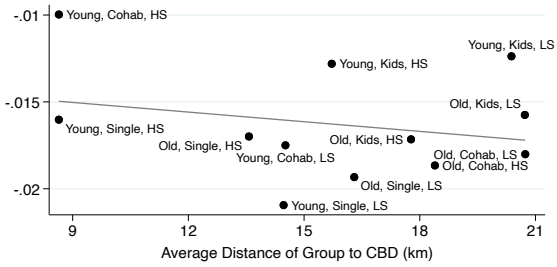
(a) Housing Expenditure Share (1- α^{of})



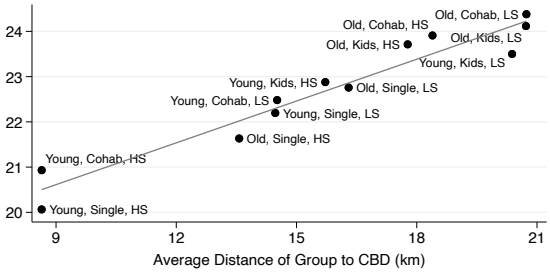
(b) Gravity coefficient, PPML IV (ν^{of})



(c) Commuting cost parameter (ϕ^{of})



(d) Amenity-weighted distance to CBD ($\mathcal{B}_i^{of}$)



Notes: Average distance of group to CBD is the residential population-weighted distance from residences to the CBD by group. Amenity-weighted distance from the CBD is the distance from a parish to the CBD, weighted by the group-specific amenity recovered from the model inversion. Larger values imply that a group prefers amenities in places further away from the CBD.

locations where the unit price of housing is generally lower. Figure 9a, however, reveals that groups with low expenditure shares tend to occupy the suburbs, suggesting that other mechanisms must be driving the spatial sorting.

In a monocentric city like Copenhagen, commuting costs act as a concentration force since utility declines the farther workers live from the CBD, where most jobs are located. All else equal, workers who dislike commuting will prefer to live nearer to the CBD, offering an alternative explanation for spatial sorting. In Figure 9b, however, we observe a negative relationship between the PPML-IV gravity parameters estimated in Table 1 and the average distance at which workers locate. Since commuting-averse groups live in the suburbs, and vice versa, heterogeneous attitudes towards commuting are also not a compelling explanation of the observed spatial sorting pattern.[29]

One of the strengths of a quantitative urban model is that it can decompose the determinants of location choice into the pecuniary elements that are canonical within economic models, such as the expected wage at the workplace, commuting costs, and housing rents, and an exogenous

[29]The gravity parameters estimated are a combination of the preference parameters ε^{of} and the commuting cost parameter ϕ^{of}, therefore one could argue that the relationship displayed in Figure 9b is driven by the preference shocks. In Figure 9c, we show the same scatterplot for the commuting cost, which also shows a negative relationship between the parameter and the group average distance to the CBD.

structural residual that captures the amenity value of a location. The amenity value of a place is likely very specific to the group. For example, parents are likely to derive utility from the presence of a nearby playground, whereas singles may prefer places with bars and coffee shops to socialize. Since neither housing nor commuting preferences can explain the observed residential sorting, it seems intuitively plausible that heterogeneous amenity values are the driving force.

Figure 9d reveals that the suburbanization of amenity preferences is a striking predictor of the suburbanization of residences. In other words, households live where they find the amenities they enjoy at the given stages of their lives. While this pattern is striking, one concern is that amenities are a structural residual of the model and could be capturing a number of factors. To shed some light on the nature of amenities for which groups have heterogeneous preferences, we compute indices of observable proxies for amenities that have been widely discussed in the literature: consumption amenities, natural amenities, and school quality. We correlate these amenity indices with the model-inverted amenity, allowing for interactions with group characteristics in Table 2. The reference group (absorbed by fixed effects) is young, low-skilled single people without children. The results show that older people, couples, and parents value consumption amenities less, while the high-skilled value them more than the reference group. For example, the quantitative estimates imply that the utility of seniors increases in the consumption amenity index at an elasticity that is two percentage points lower than that of the reference group. The pattern reverses when considering the natural amenity index, implying that groups valuing natural amenities relatively more also value consumption amenities relatively less. While parents, as expected, value school quality more than the reference group, the effect is quantitatively small, which is consistent with the generally high quality of public schools in Denmark. The combined insight from Figure 9d and Table 2 is that seniors, couples, and parents sort into suburbs because they appreciate natural amenities, while the young, single, childless, and high-skilled prefer the consumption amenities found in the city center.

5.3 Model Counterfactuals for Sorting

To assess the relative contributions of different sources of group-specific heterogeneity to the spatial sorting of groups, we employ model-based counterfactuals. Specifically, we eliminate group-specific heterogeneity in selected primitives and solve for the counterfactual spatial general equilibrium using the mapping from primitives to endogenous variables discussed in Section 4.5.

The first counterfactual assumes that all groups have the same housing expenditure share, which is set equal to the average housing expenditure share across groups, and solves for a counterfactual equilibrium under this assumption. The following counterfactuals remove the variation in the Fréchet shape parameter, the elasticity of commuting cost with respect to distance, and the variation in the labor input shares in the production function across locations. The final counterfactual

Table 2: Relative Amenity Preferences by Groups

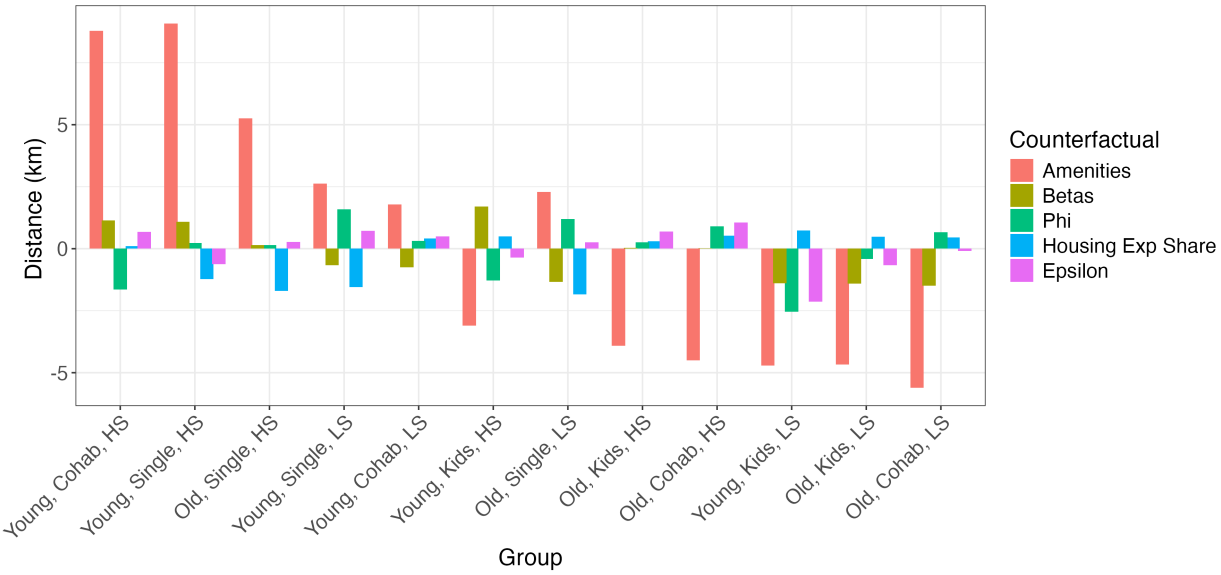
Interaction effects	Senior	Cohabiting without children	Cohabiting with Children	High-skill
Consumption amenity index	-0.0213^{***}	-0.0293^{***}	-0.0133^{***}	0.0183^{***}
	(0.0014)	(0.0017)	(0.0017)	(0.0014)
Natural amenity index	0.0332^{***}	0.0124^{***}	0.0100^{***}	-0.0099^{***}
	(0.0026)	(0.0031)	(0.0031)	(0.0026)
School quality index	0.0018^{**}	0.0013	0.0023^{**}	0.0009
	(0.0008)	(0.0010)	(0.0010)	(0.0008)
Observations		3,421		
Group FE		Yes		
Parish FE		Yes		
R-squared		0.7658		

Notes: The table shows the interaction coefficients from one regression at the parish-group level. The specification is a log-log model with the inverted model amenity by region-group as the dependent variable and amenity indices based on observed data interacted with group as independent variables, controlling for group and location fixed effects. The reference group are the young, single and low-skilled. The consumption amenity index includes the density of cinemas, restaurants, bars, and street food establishments. The natural amenity index includes the number of marinas, location on the coast, the share of water, and the share of green, and non-built-up land. School quality is based on centralized key-stage test scores in Danish and mathematics at the conclusion of primary education (9th grade). All indices are computed by taking the geometric mean across selected variables after normalizing each of the inputs by the mean. Standard errors are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

removes differences in residential amenities across groups by assuming that all groups attach the same amenity value to a location, which we set equal to the unweighted average of the estimated amenities of each group for the location in the baseline.

Figure 10 summarizes the results of these counterfactuals. It shows that eliminating variation in each parameter across groups has only a minimal impact on the average distance between a group's residential location and the city center, except in the case of amenity variation. When all groups are assumed to evaluate amenity values identically, the spatial sorting of groups changes substantially. Initially, more central groups increase their average distance to the CBD, while more suburban groups decrease it. These findings suggest again that differences in how groups perceive residential amenity values play a central role in shaping their location choices.

Figure 10: Parameter Counterfactuals - Change in Average Distance to the CBD



Notes: The figure plots the relative importance of the five model parameters – amenities, labor demand (β_i^o), commuting cost elasticity (ϕ^{of}), housing expenditure shares (α^{of}), and idiosyncratic preferences (ε^{of}) – in determining the average distance of a group's residential location from the city center. The x-axis order the twelve demographic groups by the most central to the least central, while the y-axis plots the change in the average distance in kilometers to the CBD.

6 The Effect of Demographic Change on Cities

The stark differences in location choices across different groups in our data suggest that demographic changes such as population aging or changing fertility could have first order impacts on the internal organization of cities. In this section, we use our estimated model to explore the effects of three prominent demographic changes – population aging, reductions in fertility and increases in the share of single households.

As discussed in Section 2, Copenhagen experienced only marginal demographic changes along these three dimensions during our sample period. To estimate the effects of plausible levels of demographic changes in cities, we use published census data for Tokyo, a case known for its dramatic demographic aging.[30] Table 3 shows how the adult population of the Greater Tokyo area changed between 1980 and 2020. The share of adults living with children (under the age of 18) more than halved from 43% in 1980 to just 19% of the adult population in 2020. Similarly, we see an increase in singles from 32% to 40% of the adult population. Finally, the share of the population

[30] OECD (2015) show in their Figure 1.20 that Tokyo has the highest annual growth rate of the 65+ population among the major OECD metropolitan areas over the 2001-2011 period, with the annual growth rate of the over 65 population exceeding 4% compared to a growth rate of the total population of the city of less than 1%.

which is 45-64 increases from 31% to 46% of the 18-64 population.[31]

Table 3: Demographic Change in the Greater Tokyo Area

	Age		Couples		Parents	
Year	18-44	45-64	Single	Couple	Parents	Non-Parents
1980	69%	31%	32%	68%	43%	57%
2020	54%	46%	40%	60%	19%	81%

Notes: The table shows percentage shares of different types of people in the total adult population (those 18 or older) in the Greater Tokyo Area of either the 1980 or 2020 Japanese census. The sample for the age shares covers all individuals aged 18-64 in the Greater Tokyo Area in 1980 and 2020, while the sample for the couples and parents shares covers all individuals aged 18 or older in the Greater Tokyo Area in 1980 and 2020. The data comes from different tabulations of the Japanese population census results as described in detail in Online Appendix F.

To explore the effect of such dramatic demographic change on the internal organization of cities, we simulate a demographic shift comparable to the one observed in Tokyo between 1980 and 2020 using our estimated model in Copenhagen. Specifically, we solve for three model counterfactuals for Copenhagen – an increase in the share of singles, a decrease in the number of couples with children, and an increase in the share of seniors by the same percentage point change observed in Tokyo between 1980 and 2020. To isolate the effects of these three demographic changes, we hold the composition of the population across other demographic dimensions constant. For example, to increase the share of old people we reallocate individuals from the young to the older population within the same skill, cohabitation and fertility category. We do so equally across all skill, cohabitation and fertility subgroups of the population as the published Japanese census data does not show how aging differs across subgroups. We apply the same logic for the other two counterfactuals simulating a decrease in the number of couples with children and an increase in the share of seniors. We undertake these counterfactuals using the same assumptions as in Section 5, i.e. we consider a closed city with fixed productivity and amenities and hold the supply of floor space and its allocation to commercial and residential use constant.[32]

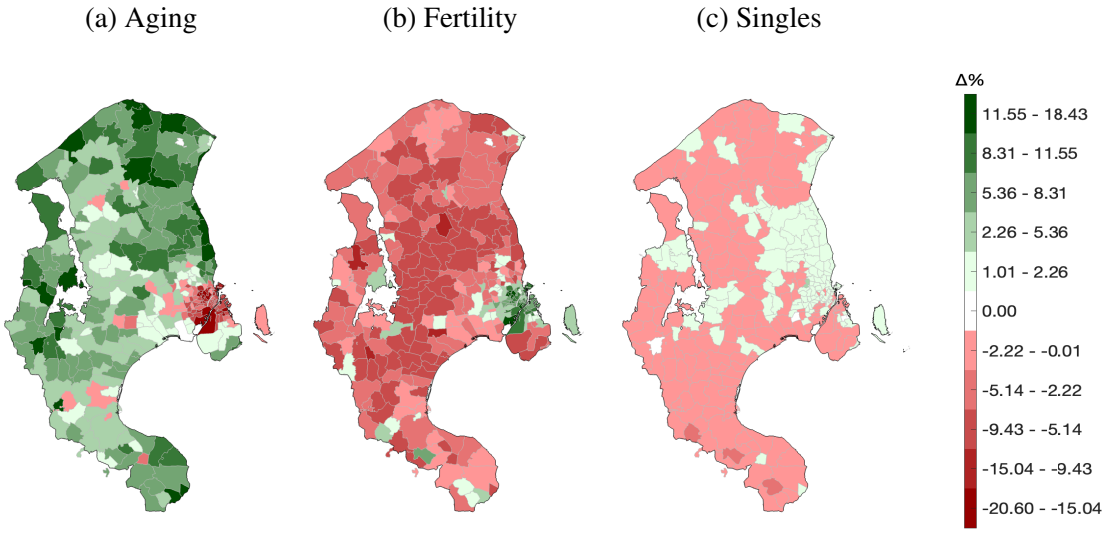
Figure 11 shows the results of these model counterfactuals for each of the three demographic changes on the total residential population of each of the parishes in Copenhagen. The figure shows several striking patterns for how the population in Copenhagen will relocate in response to demographic changes. Figure 11a shows that population aging is a substantial decentralizing force. Residential demand shifts away from central locations and into more peripheral and suburban areas

[31] In 2011, which is the baseline year of our model, the comparable shares for Copenhagen are 25% of the total adult population live with children under the age of 18, 41% are single, 60% are between 18 and 44 and 40% between 45 and 64. The shares in our model are slightly different as the baseline model only considers those who are working.

[32] Holding the amenities and productivity of each location fixed means that all the changes we measure are likely to be a lower bound of the results if we allowed these to be at least in part endogenous to the distribution of the population.

consistent with the lower amenity value that older people attach to central parts of the city.

Figure 11: Counterfactual Changes in the Residential Population



Notes: This figure shows the results of the model counterfactuals for each of the three demographic changes on the total residential population. Each panel maps the percentage change in residential population (Δ%) under a specific counterfactual scenario, holding other factors constant. Panel (a) shows the predicted change due to aging alone. Panel (b) shows the effect of the decrease in fertility. Panel (c) shows the impact of the increase in the share of singles. Color gradients represent the magnitude of percentage change, with darker shades indicating stronger effects.

Figure 11b reveals that a reduction in the number adults with children would have the opposite effect. Consistent with both the event study estimates and the model inferred amenities, groups with children prefer more suburban locations. Fewer children therefore increase demand for more central locations. Figure 11c shows that the relatively small increase in the number of singles observed in Tokyo has a modest centralizing effect on the residential population. This is in line with the event study estimates that show that cohabitation is correlated with more suburban location choices and the same pattern is also visible in the inverted model amenities. A higher share of the population being single therefore also results in an increased demand for central locations.

Appendix Section F shows maps for the counterfactual impacts of these three demographic changes for employment (Appendix Figure F1), residential floor space prices (Appendix Figure F2), and commercial floor space prices (Appendix Figure F3). These maps show that demographic change also affects the distribution of employment, but the effects are more muted. This is consistent with the reduced form evidence showing that different groups sort more in terms of residential location than workplace location. The effect of demographic change on residential floor space prices largely follows the distribution of residential population. It is not entirely proportional to the population changes, due to the different expenditure shares and income levels across groups. A

notable feature of the effect of an increase in the number of singles on residential floor space prices is that it not only triggers relative changes across locations but also lifts population weighted average prices for the entire city. This is consistent with the reduced form findings that cohabitation and separation have substantial effects on the amount of floor space consumed and implies that an increase in the number of singles could be a substantial driver of house prices in cities in the future.

While our model counterfactuals isolate the effects of different demographic changes, the most likely future scenario is one in which we observe all of these demographic changes at the same time. The experience of Tokyo over the period from 1980 to 2020 suggests that population aging goes hand in hand with fewer children and more singles. If this is the pattern of future demographic changes, our results suggest that the decline in fertility and increase in singles will work against the effect of population aging in terms of their effect in decentralizing economic activity in cities. The combined effect of demographic change on cities is therefore likely much more muted than the effect of each of these demographic changes on its own, which suggests that cities could be quite resilient in the face of demographic change. Consistent with this Appendix Figure F4 shows the results for a counterfactual that combines all the three demographic changes observed in Tokyo between 1980 and 2020. The results show that this particular combination of demographic change has a moderately decentralizing effect on population and employment, but the effects are much more muted than the effect of aging or fertility changes on their own.[33]

7 Conclusion

Our findings highlight that intra-urban location choices are fundamentally shaped by how different demographic groups value amenities. These preferences vary systematically across age groups and family types. Given that amenities are unevenly distributed within cities, demographic megatrends - such as population aging, declining fertility, and rising singlehood - have the potential to significantly shift spatial demand patterns.

Such shifts in demand caused by demographic change present potential challenges for urban infrastructure and housing markets, as changes in population composition can alter the demand for central versus peripheral locations. However, our counterfactual analysis offers a more nuanced perspective. We show that the spatial implications of these demographic forces often counteract one another. For example, while declining fertility and increased singlehood tend to intensify demand for central urban locations, aging populations exert decentralizing pressures. This interaction

[33] One limitation of our demographic counterfactuals is that we undertake them in our baseline model, in which we only consider the working population. Cities obviously also contain substantial amounts of non-working residents such as pensioners. Furthermore, the experience of Tokyo has seen rapid increases in the number of pensioners. In future work we will extend the model to include also non-working groups to explore the sensitivity of the demographic results to the inclusion of non-working population in the model.

suggests a degree of resilience in urban spatial structures to demographic change.

We also find that mobility considerably decreases with age and there is little downsizing of floor space consumption with age, even once individuals become pensioners and are no longer attached to the labor market. Our quantitative exercise focuses mostly on individuals attached to the labor market, and we leave for future research to understand why location choices and floor space consumption change so little after big life changes in older age. The results also indicate there is a dynamic component to the structure of cities, especially if population changes and the city grows.

References

Adda, J., Dustmann, C., and Stevens, K. (2017). The Career Costs of Children. *Journal of Political Economy*, 125(2):293–337.

Ahlfeldt, G. M., Heblich, S., and Seidel, T. (2021). Micro-Geographic Property Price and Rent Indices. *CEP Discussion Paper*, 1782.

Ahlfeldt, G. M., Heblich, S., and Seidel, T. (2023). Micro-Geographic Property Price and Rent Indices. *Regional Science and Urban Economics*, 98:103836.

Ahlfeldt, G. M., Redding, S. J., Sturm, D. M., and Wolf, N. (2015). The Economics of Density: Evidence from the Berlin Wall. *Econometrica*, (6):2127–2189.

Albouy, D. and Faberman, J. R. (2025). Skills, Migration, and Urban Amenities over the Life Cycle. NBER Working Paper 33552, National Bureau of Economic Research.

Allen, T., Arkolakis, C., and Li, X. (2015). Optimal City Structure. Working paper.

Ang, A., Angel, D., and Parkhomenko, A. (2024). Amenities in quantitative spatial models. *Working Paper*.

Anstreicher, G. and Venator, J. (2024). To Grandmother's House We Go: Informal Childcare and Female Labor Mobility. Working paper.

Badilla Maroto, M., Faber, B., Levy, A., and Munoz, M. (2024). Senior Migration, Local Economic Development and Spatial Inequality. Working paper.

Bayer, P., Ferreira, F., and McMillan, R. (2007). A Unified Framework for Measuring Preferences for Schools and Neighborhoods. *Journal of Political Economy*, 115(4):588–638.

Bayer, P. and Timmins, C. (2007). Estimating equilibrium models of sorting across locations. *Economic Journal*, 117(518):353–374.

Ben-Porath, Y. (1967). The Production of Human Capital and the Life Cycle of Earnings. *Journal of Political Economy*, 75(4):352–365.

Berry, S. (1994). Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, 25(2):242–262.

Bonnet, C., Gobillon, L., and Laferrère, A. (2010). The Effect of Widowhood on Housing and Location Choices. *Journal of Housing Economics*, 19(2):94–108.

Bordeu, O. (2024). Commuting Infrastructure in Fragmented Cities. Working paper.

Borusyak, K., Jaravel, X., and Spiess, J. (2024). Revisiting Event-Study Designs: Robust and Efficient Estimation. *Review of Economic Studies*, 91(6):3252–3285.

Browning, M. and Ejrnæs, M. (2009). Consumption and Children. *The Review of Economics and Statistics*, 91(1):93–111.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics*, 225(2):200–230.

Coeurdacier, N., Combes, P.-P., Gobillon, L., and Oswald, F. (2024). Fertility, Housing Costs and City Growth. Working paper.

Cortés, P. and Pan, J. (2023). Children and the Remaining Gender Gaps in the Labor Market. *Journal of Economic Literature*, 61(4):1359–1409.

Couture, V., Gaubert, C., Handbury, J., and Hurst, E. (2024). Income Growth and the Distributional Effects of Urban Spatial Sorting. *Review of Economic Studies*, 91(2):858–898.

Fernández, R. and Wong, J. C. (2014). Divorce Risk, Wages and Working wives: A Quantitative Life-Cycle Analysis of Female Labour Force Participation. *The Economic Journal*, 124(576):319–358.

Foerster, H. (2024). Untying the Knot: How Child Support and Alimony Affect Couples' Dynamic Decisions and Welfare. *Review of Economic Studies*, rdae105.

Gaubert, C. and Robert-Nicoud, F. (2024). Sorting to Expensive Cities. Working paper.

Gautier, P. A., Svarer, M., and Teulings, C. N. (2010). Marriage and the City: Search Frictions and Sorting of Singles. *Journal of Urban Economics*, 67(2):206–218.

Heblich, S., Redding, S. J., and Sturm, D. M. (2020). The Making of the Modern Metropolis: Evidence from London. *The Quarterly Journal of Economics*, 135(4):2059–2133.

Heckman, J. J. (1976). A Life-Cycle Model of Earnings, Learning, and Consumption. *Journal of Political Economy*, 84(4, Part 2):S9–S44.

Jain, V. and Arai, Y. (2019). Case study on tokyo metropolitan region, japan. Policy Paper Series 3, World Bank Group, Washington, D.C.

Kleven, H., Landais, C., and Leite-Mariante, G. (2024). The Child Penalty Atlas. *Review of Economic Studies*, page rdae104.

Kleven, H., Landais, C., and Søgaard, J. E. (2019). Children and Gender Inequality: Evidence from Denmark. *American Economic Journal: Applied Economics*, 11(4):181–209.

Komissarova, K. (2022). Location Choices over the Life Cycle: The Role of Relocation for Retirement. Working paper.

Kuminoff, N. V., Smith, V. K., and Timmins, C. (2013). The new economics of equilibrium sorting and policy evaluation using housing markets. *Journal of Economic Literature*, 51(4):1007–1062.

Liu, L., Wang, Y., and Xu, Y. (2024). A Practical Guide to Counterfactual Estimators for Causal Inference with Time-Series Cross-Sectional Data. *American Journal of Political Science*, 68(1):160–176.

Lucas, R. E. and Rossi-Hansberg, E. (2002). On the Internal Structure of Cities. *Econometrica*, 70(4):1445–1476.

Meghir, C. and Pistaferri, L. (2010). Earnings, Consumption and Life Cycle Choices. In Ashenfelter, O. C. and Card, D. C., editors, *Handbook of Labor Economics*, volume 4B, pages 773–854. Elsevier, Amsterdam.

Mincer, J. A. (1974). *Schooling, Experience, and Earnings*. National Bureau of Economic Research (NBER).

Modigliani, F. and Brumberg, R. H. (1954). Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data. In Kurihara, K. K., editor, *Post Keynesian Economics*, page 388–436. Rutgers University Press, New Brunswick, NJ.

Moreno-Maldonado, A. and Santamaria, C. (2024). Delayed Childbearing and Urban Revival. Working paper.

OECD (2015). *Ageing in Cities*. OECD Publishing, Paris.

Owens III, R., Rossi-Hansberg, E., and Sarte, P.-D. (2020). Rethinking Detroit. *American Economic Journal: Economic Policy*, 12(2):258–305.

Redding, S. J. (2024). Quantitative Urban Economics. NBER Working Paper 33130, National Bureau of Economic Research.

Redding, S. J. and Sturm, D. M. (2024). Neighborhood Effects: Evidence from Wartime Destruction in London. NBER Working Paper 32333, National Bureau of Economic Research.

Rich, J. and Overgaard Hansen, C. (2016). The Danish National Passenger Model – Model Specification and Results. *European Journal of Transport and Infrastructure Research*, 16(4).

Sun, L. and Abraham, S. (2021). Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects. *Journal of Econometrics*, 225(2):175–199.

Tsivanidis, N. (2023). Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogota's TransMilenio. Working paper.

Weiwu, L. (2024). Unequal Access: Racial Segregation and the Distributional Impacts of Interstate Highways in Cities. Working paper.

Wooldridge, J. M. (2015). Control Function Methods in Applied Econometrics. *Journal of Human Resources*, 50(2):420–445.

A Data

In this Appendix Section, we provide an overview of how our datasets were constructed, and describe the key variables used in our analysis. Our main dataset is an annual panel of all people living in Denmark over the period of 1986 to 2019.

A.1 Building the population panel

Data on the population comes from a series of high-quality statistical registers maintained by the Danish Ministry of the Interior and accessed through Statistics Denmark. All people living in Denmark for more than three months are legally required to be registered and obtain a unique ID called *Det Centrale Personregister* (CPR). CPR registration is necessary for accessing healthcare, legal employment, personal finance, utilities, education and housing among others and, as a result, non-compliance is close to null.[34] Similarly, errors and delays do occur among the general population when moving addresses or changing jobs.

The basic information from CPR contains residence, workplace, birthday and various other social characteristics such as biological or adopted children, spouses and/or cohabiting partner, and education achieved. Importantly, it also contains a unique time-invariant ID for each registered person. We can therefore link the CPR register with data from the tax authorities on wage income, capital income, transfers from government and unemployment insurance, taxes, mortgage and assets owned, imputed cost of owner-occupied housing and various other outcomes of interest. We build a yearly panel, since some of the datasets update the information at the start of each year, in January.

Furthermore, each dwelling has a unique time-invariant ID. We can therefore link each citizen's residence with floor space, housing type, construction year, transaction price and assessed value for taxation purposes. We obtain data on property values from two data sources because transaction data is only available after 1994. To obtain property prices for the period before that, we use data with assessed property values. The assessments data have been consistently criticised for their low accuracy and, as a result, a political decision was taken to freeze them at their 2012 levels: after that, we do not have any information at all.

More specifically, to build our final dataset we use several data sources from Statistics Denmark. We use the register BEF to get family identifiers and addresses, and the register UDDA to obtain the highest level of education achieved by the individuals[35]. For work characteristics, we use the register IND to obtain information on income, IDAP to record years of work experience,

[34] Illegal immigrants are not registered in CPR and emigrants have been known to occasionally forget to inform the state of their departure from the country; both factors are believed to cause a discrepancy between CPR registration and actual residential population by up to 0.2%.

[35] We only observe this information for individuals who completed their education in Denmark.

IDAN for data on workers wages, occupation and sector, IDAN for data on workplace ID and industry, and RAS for worker's full- or part-time status. We use the death census (DOD2019) to flag if individuals have died, and create a child census to determine the number of children individuals have.

Our population sample consists of all adult individuals, with at least 18 years old, living inside the Copenhagen Metropolitan Area. For analyses that involve travelling to work, we exclude those out of the labor force, self-employed, workers working in the same hectare cell as they live, and people above the age of 65. Above this age, there is substantial selection bias with respect to remaining on the labour market.

A.2 Geographic Scope: Copenhagen Metropolitan Area

We exclude people that live or work outside of the Copenhagen Metropolitan Area (*HT-området*). Until 2007, the area was used as the geographical scope of public transport coordination in the capital area; apart from that, it has and has had no significance in the Danish political economy and border discontinuities are unlikely to contaminate our results. This is a large area that captures the overwhelming majority of commuter flow into the CBD. At the same time, it is sufficiently small to be almost perfectly monocentric: only small, historic cities in Roskilde and Elsinore break the pattern of orientation towards the Copenhagen CBD.

The only change made to the Copenhagen Metropolitan Area is that we systematically exclude the small, almost unpopulated island of Hesselø located 30 kilometers from Zealand.

A.3 Computation of travel times

Driving travel times are time-invariant and are based on the road network of late 2019. This includes road congestion as measured by Google Maps. For public transport, we use the network as it was in late 2019 with the exception of the newly opened M3 metro line, which would not have been open at any time during our population panel. We include a station waiting penalty that depends on the rush-hour headway between buses and trains. We assume average speeds of 16 km/h for buses, 34 km/h for metro and 42 km/h for suburban and regional trains (*S-tog* and *regionaltog*, respectively). For cycling and walking, the length of the shortest footpath was found. It was then assumed that walking occurred at 5 km/h and cycling at 15.54 km/h.

We use the imputed travel mode shares from the Danish Traffic Model (*Landstrafikmodellen, LTM*, which is calculated based on a set of zones that are larger than the four-hectare cells that are the basis of the travel time calculations. Between each pair of four-hectare cells, we calculate an aggregate travel time that is the simple average of travel times by driving, public transport, walking and cycling weighted by the respective mode shares. This entails that pairs of cells located close to

each other in the city centre can have relatively high transport times associated with them because a large proportion of travellers will walk or cycle.

The benefit of using mode-weighted travel times is that we address potential issues with endogeneity in the simultaneous choice of transport mode and residential-employment pair. As a simple robustness check, we replicate all key findings using straight-line distance instead.

A.4 Variable definition

A.4.1 Outcomes

Our seven outcomes are travel time from residence to work, from residence to CBD, from workplace to CBD, straight line distance from residence to work, from residence to CBD, from workplace to CBD, and floor space per adult. Travel time variables are all in minutes, straight line distance variables are in kilometers, and floor space per adult is in square meters. All of the travel time and straight line distance variables are taken directly from the travel time matrices described in the previous section. Floor space per adult is defined as the total floor space of the household/adddress/dwelling divided by the number of floor space-consuming adults. We define a floor space consuming adult as anybody who is 18 or older that does not live with their parents. We also winsorize the floor space at the top and bottom 1 percent.

A.4.2 Mobility variables

Denmark is split into 29 Commuting Zones (or Travel to Work Areas, called Pendlingsområder in Danish) which are constructed from municipalities, and the largest of which is our study area, the Copenhagen Metropolitan Area (CMA). To measure mobility within the CMA and between the rest of Denmark, we first construct a measure of residence and workplace movement by tracking changes in residence and workplace grid cells for each individual over time. Since the municipalities boundaries never cross the Commuting Zone boundaries, each grid cell is assigned a unique Commuting Zone.

Using the complete register of adults in Denmark, we are able to track movements into and out of Copenhagen and not only movements within our study area.

If an individual disappears from our study area before 2019 (our final year of data), this means they have either moved out of the country or died. Using the DOD register, we are able to identify which disappearances are from death and which are from emigration and thus we can also track mobility in and out of Denmark.

When measuring workplace mobility, if an individual is missing a workplace grid cell this indicates that they were unemployed on January 1st. In our measure of workplace mobility, we do not consider transitions in and out of employment to be a workplace move. Additionally, if an

individual becomes unemployed for a period and then re-starts employment at the same previous workplace then we do not count this as a move either.

A.4.3 Life events

We study eleven main life events that we separate into early life events and late life events. The early life events include the birth of the first three children, the beginning of the first three cohabitations, and the end of the first two cohabitations. The late life events are becoming a pensioner, becoming a widow, and becoming an empty-nester. The decision to split the life events into two separate groups for the event study analysis was determined by the median age of treatment, which was younger than 45 for all of the early life events and older than 45 for all of the late life events. The decision of which events to include in our analysis was determined by the frequency of the events in our main dataset, and we chose to only include events which were experienced by at least 2.5 percent of the population.

We define the beginning of a cohabitation as the year that we observe a change in cohabitation status of an individual.[36] Since the registry data is recorded on January 1st of each year, this means that we observe the change in cohabitation status at most 12 months after the change occurs. Similarly, we determine a separation when a cohabitation ends, that is, when a couple stops living together.[37]

For childbirth, we consider the time of treatment to be the conception/pregnancy rather than the moment of childbirth. This is because expecting parents have nine months of pregnancy to prepare for the arrival of their child and therefore will update their preferences long before the child is born. Because of this, unlike with the cohabitation variable, we define the time of treatment as the year before we observe the birth of a child.

Widowhood is the year after the death of a cohabitation partner. Pension is the year the main source of income changes to pensions for the last time. The event of becoming an emptynester is defined as the year when all of a parent's children have turned 18 and moved out of the parent's home for the final time.

[36] Cohabitation is determined by when a couple starts living together. It will include married people, registered partnerships, and different-sex couples that live together, with no other adults present, no biological connection and with an age difference of less than 15 years.

[37] If someone separates and within less than a year they are already cohabiting with another person, we do not consider that a separation. Additionally, we do not consider a separation if the partner has died.

B An Econometric Model of Life Cycle in the City

In this Appendix Section, we provide details on the methodology used to estimate the event-studies and the necessary assumptions for the estimation.

The results of the Subsection 3.2 suggest that location choices and floor space consumption have clear life cycle patterns. However, life events, such as marriage and having children, could also play an important role in driving these trends we observe. In this subsection, we derive an econometric model aiming to disentangle what is driven by aging alone and what are the effects of a series of life events.

Our econometric framework supports our analysis and aims to answer the following questions: do life events impact the location choices and floor space consumption of individuals? Which life events are more important at explaining the changes in these outcomes? How important are the life events relative to aging alone? Below, we lay down a simple econometric model that guides our approach to estimating the effects of life events on a series of economic outcomes that we are interested in. We follow closely the framework proposed by Liu et al. (2024) and Borusyak et al. (2024), adapting it to decompose the total effect of life events into the leads and lags of separate life events, such as cohabiting, having children, etc.

Let $Y_{i,a}$ denote the distance from residence to the CBD– or any other outcome of interest– for person i at age a. Denote $\mathbf{D}_{i,a}$ as a vector of treatment status of the L life events. We model $Y_{i,a}$ to be determined by

$$Y_{i,a} = \theta_i + \eta_a + \mathbf{D}'_{i,a}\tau + \varepsilon_{i,a} \tag{B.1}$$

such that it depends on an individual specific factor captured by θ_i, an age specific factor captured by η_a, and $\mathbf{D}_{i,a}$ is a vector where each row indicates if the individual i has been treated at age a with life event $l \in \{1,...,L\}$. Therefore, the vector τ contains the treatment effects of the life events. Let $Y_{i,a}(0)$ be the outcome of individual i and age a when $\mathbf{D}_{i,a} = 0$, that is, when the individual has not been treated with any life event.

Assumption 1 (counterfactual) The potential outcome of individual i at age a untreated with any event is known and equal to $Y_{i,a}(0) = \theta_i + \eta_a + \varepsilon_{i,a}$. This assumption rules out any form of spatial or temporal interference, such as anticipation effects for any of the events.

Assumption 2 (additive separability) The parametric function of the effects of life events on outcome $Y_{i,a}$ is known and requires additive separability, such that $\mathbf{D}'_{i,a}\tau = \sum_{\ell=1}^{L} d^{\ell}_{i,a}\tau^{\ell}$, where $d^{\ell}_{i,a}$ indicates if individual i with age a has been treated with life event ℓ and τ^{ℓ} is that life event treatment effect.

To develop intuition over what are average treatment effects we are after, we consider the potential outcomes framework for L life events. Let us define Z as a set of vectors containing all possibilities of treatment and control for each life event. If there are L events, the number of vectors inside Z equals $z = 2^L - 1$ (at least one entry equals one), and we estimate treatment effects for each of them. We can define the following potential outcomes:

$$\delta_{i,a}^{z} = Y_{i,a}(Z = z) - Y_{i,a}(0)$$

where δ^z is a linear combination of the average treatment effects. To exemplify, let us consider the case where we have 2 life events so $z = 3$. Under additive separability, the relevant average treatment effects are estimated as:

$$\begin{aligned}
\delta_{i,a}^{z_1} &= Y_{i,a}(1,0) - Y_{i,a}(0) = \tau^1 \\
\delta_{i,a}^{z_2} &= Y_{i,a}(0,1) - Y_{i,a}(0) = \tau^2 \\
\delta_{i,a}^{z_3} &= Y_{i,a}(1,1) - Y_{i,a}(0) = \tau^1 + \tau^2
\end{aligned}$$

Assumption 3 (strict exogeneity) $\varepsilon_{i,a} \perp\!\!\!\perp \{\theta_j, \eta_s, \mathbf{D}_{j,s}\} \quad \forall i, j \in \{1,...,N\}$ and $a, s \in \{1,...,A\}$.

Under Assumptions 1-3, we are able to estimate our model using the following steps:

Step 1: Use the sample of untreated to calculate the individual and age fixed effects.

Step 2: Predict the counterfactual and obtain the individualistic $\hat{\delta}_{i,a}^{z}$.

Step 3: Decompose the estimated into the sum of treatment effects of life events:

$$\hat{\delta}_{i,a}^{z} = \sum_{\ell=1}^{L} d_{i,a}^{\ell} \tau^{\ell} + u_{i,a}$$

Assumption 3 is necessary for interpreting the effects of life events as causal. Life events are not randomly allocated, and we don't want to claim that our estimates represent the causal effects of life events on location choices and floor space consumption, as there are other things that may be changing when people choose to go through these life changing decisions, such as a first cohabitation or a first child.[38] Nonetheless, we show (descriptive) evidence that when people choose to experience life events, their location within the city changes, and these results are what matters for our analysis.

[38] First cohabitation is the first life event experienced by close to 75% of our estimating sample.

C Stylized Facts

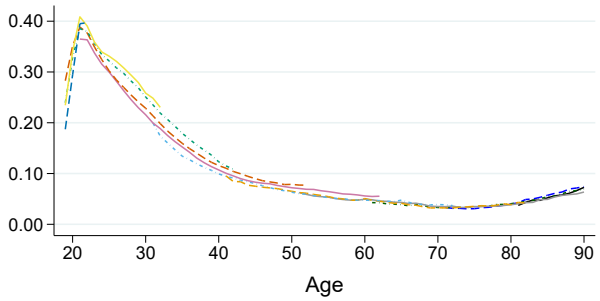
In this Appendix Section, we provide theoretical and empirical content supporting the exhibits in Section 3. For the empirical results, we consider the same methodology presented in the main paper, and provide additional heterogeneities and robustness.

C.1 Mobility over the Life Cycle

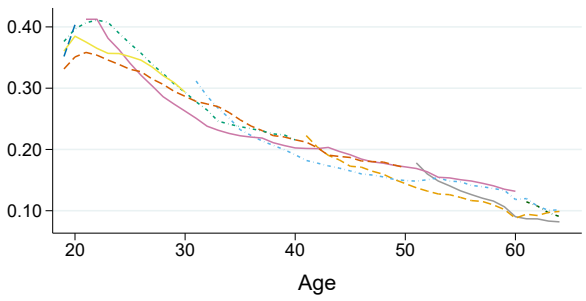
In Subsection 3.1, we show that mobility is higher in the earlier years of adult life, both in terms of residence and in terms of workplace. Our outcomes are presented for the whole CMA population, and pooling together all the years of our dataset. The main concern one may have is that our measure is combining both age and cohort effects, and if the latter are strong enough, we could see different patterns for each cohort. Figure C1 replicates the methodology splitting the sample by birth-decade cohort. Each line in the graph represents the mobility probability for a given cohort, and they don't always cover the entire period because we can only observe the same cohort for at most 40 years. The figure suggests that cohort effects are not as important as the age effects in explaining the pattern observed in Figure 2. The peak in the probability of moving either residence or workplace in the early twenties is still present for all cohorts. Figure C1a shows a stable life cycle pattern in residential mobility, with close to no differences across cohorts. Figure Figure C1b shows that there is a consistent pattern for workplace mobility over the life cycle, but it suggests there are slight differences in levels across cohorts.

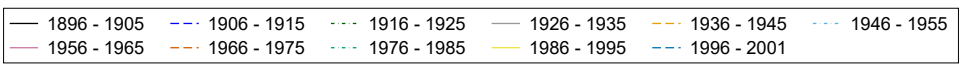
Figure C1: Residence and workplace mobility over the life cycle

(a) Probability of moving residence by cohort



(b) Probability of moving workplace by cohort





Notes: This figure shows residence and workplace mobility patterns over the life cycle by birth cohort. Panel (a) plots the probability of moving residence against age. Panel (b) graphs the probability of moving workplace against age.

In Figure C2, we replicate our methodology for a subset of commuting zones in Denmark.[39]

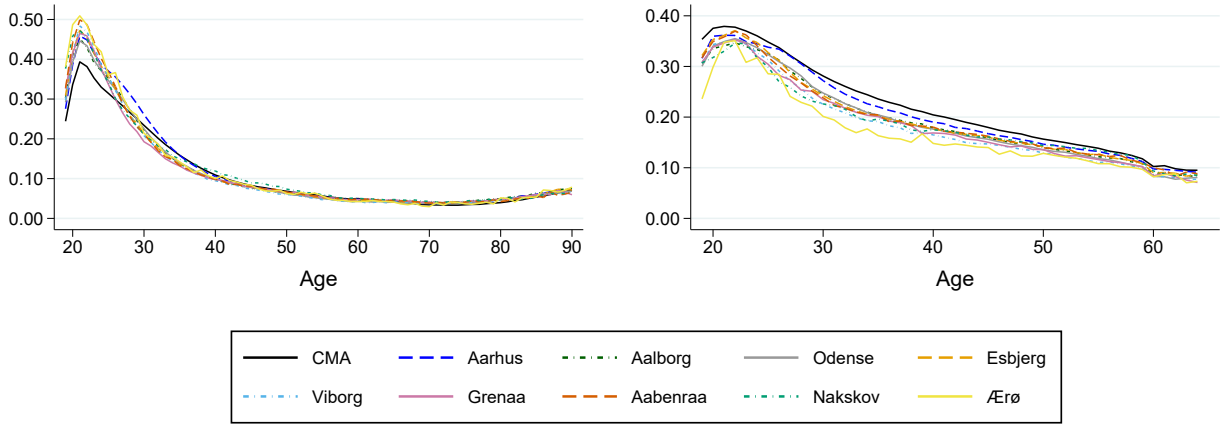
[39]There are 29 commuting zones in Denmark, and we consider 10 of these of varying population size.

The mobility pattern is observed not only in the CMA but across all regions in Denmark. Workplace mobility is overall higher in the CMA compared to other regions of the country and residence mobility is lower at younger ages, but it seems to be very similar across regions after the age of 40.

Figure C2: Residence and workplace mobility over the life cycle by commuting zone

(a) Probability of moving residence by commuting zone

(b) Probability of moving workplace by commuting zone



Notes: This figure shows residence and workplace mobility patterns over the life cycle by commuting zone. Panel (a) plots the probability of moving residence against age. Panel (b) graphs the probability of moving workplace against age.

While the previous graphs show the results for any mobility across 100m cells, our model considers parishes as spatial units. We follow by showing in Figure C3 that the pattern we find in Figure 2 is very similar. Figure C3a shows a similar pattern but with overall lower levels of mobility across parishes. Interestingly, Figure C3b is very similar in both levels and pattern.

In the previous exercises, we show that mobility has a clear age pattern, and it's much higher in early life than later in life. This is consistent across locations, across cohorts, and with different spatial units. In Figure C4, we now show the frequency of moves by distance in kilometers. Figure C4a shows that residential moves are much more likely to occur within small distances. Figure C4b shows that workplace moves are not as local as residential, exhibiting a less steep decay with distance.

C.2 Location Choices over the Life Cycle

We follow by considering heterogeneities related to the location choices over the life cycle. In Figure C5, we explore if gender and parenthood could be important factors in explaining the patterns observed in Figure 3. In this exercise, we show the mean distances between residence and

Figure C3: Residence and Workplace Mobility over the Life Cycle: Parish moves

(a) Probability of moving residence or work across parishes



(b) Conditional probability of moving within CMA



Notes: This figure shows residence and workplace mobility patterns across parishes. Panel (a) plots the probability of moving residence and workplace across parishes against age. Panel (b) graphs the conditional probability of moving residence and workplace within the CMA against age.

workplace to the CBD by age, conditional on individual fixed effects. We split our sample by gender and by ever-parenthood, that is, if individuals will ever become parents or not. We present the mean conditional distances for men in the solid black line, and for women in the dashed green line.

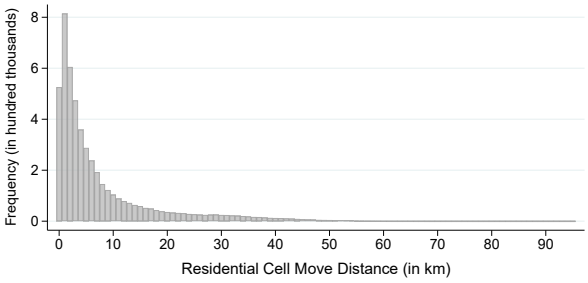
Comparing Figures C5a and C5b, two interesting patterns about residence location emerge. The first is that people that never become parents are consistently living closer to the CBD than people that become parents at some point in their lives. The difference starts increasing when individuals are in their mid-twenties and are more likely to start their reproductive life. The second pattern is that never mothers live more centrally than the never fathers, on average. The gender differences in distance from residence to the CBD are much smaller, and almost negligible for individuals that experience parenthood.

Figures C5c and C5d show the mean distance between workplace and CBD, conditional on individual fixed effects. We observe that the u-shaped pattern of workplace distance to CBD for men is independent of fatherhood status. Early in their adult life, men start working in the suburbs, centralize employment in their twenties, and start decentralizing slowly as they age. Conversely, the relative workplace location for women is both quite different from men's workplace location and quite different by motherhood status. Never mothers work much more centrally than men and other women. For women that experience motherhood, we observe a much stronger pattern of decentralization of the workplace in childbearing age.

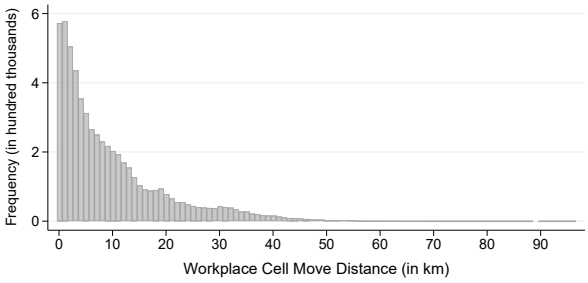
In Figure C6, we compute the same outcomes as in Figure 3 conditional on individual fixed effects, by birth decade cohort. The results show striking patterns. Figure C6a shows that there

Figure C4: Distance of Residence and Workplace Moves

(a) Distance of residence 100m cell moves (km)



(b) Distance of workplace 100m cell moves (km)

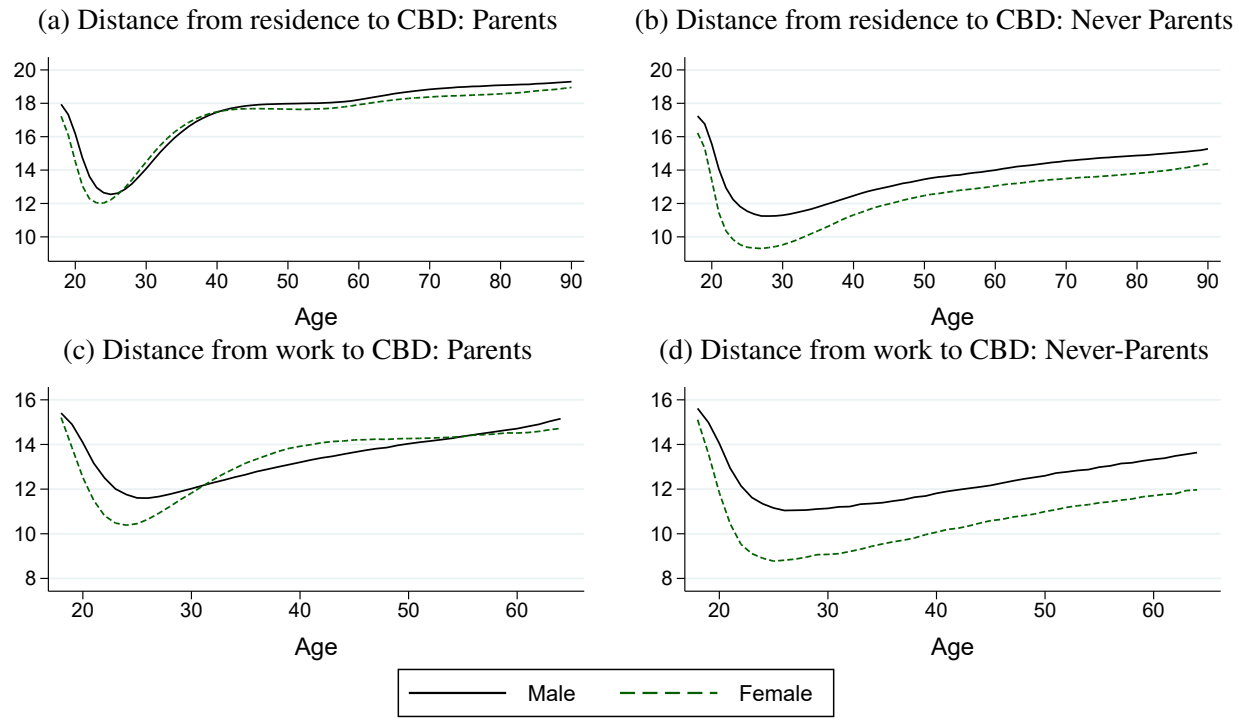


Notes: This figure plots the distribution of the kilometer-measured distance of residence and workplace relocations, in Panels (a) and (b) respectively.

is significant variation in the distance from residence to the CBD by cohort for a given age. In general, older cohorts tend to live more centrally than younger cohort at an older age. This could be explained by urban sprawl of the CMA. When older generations were young, they lived more centrally because the city was smaller, and they didn't relocate with as they aged, consistent with Figure C1. The degree of centralization that occurs in young adulthood seems to vary considerably by cohort as well.

In Figure C6b, we also observe differences in the workplace location by cohort although less pronounced than the residence location. The largest differences by cohort seem to occur for the oldest cohorts, which work more centrally after the age of 50. Figure C6d shows that the consumption of floor space per adult has a consistent pattern for cohorts, with bigger differences also occurring at later ages for the oldest cohorts. This could be related to them living relatively more central, where floor space is more expensive. Consistent with older generations having both residence and more central workplaces, C6c shows that commuting time is smaller for older generations, but the overall pattern is similar with the total population. In Figure C7, we further explore how residence and workplace location decisions vary with age by skill level. Figure C7a demonstrates that high-skilled individuals live closer to the city center than low-skilled individuals throughout their life cycle, particularly at younger ages. Figure C7b shows a similar pattern for workplace location, with high-skilled individuals consistently working closer to the city center. Figure C7c indicates that commuting times slightly increase for both groups until around age 40, after which commuting times remain relatively stable, with high-skilled workers generally experiencing somewhat longer commuting durations than low-skilled workers. Lastly, Figure C7d reveals that high-skilled individuals consume significantly more residential floor space per adult compared to their low-skilled counterparts across almost all ages, with the gap widening notably after age 40.

Figure C5: Residence and workplace location choices over age, by gender and parenthood



Notes: This figure shows residence and workplace location decisions over the life cycle by gender and parenthood status. Panels (a) and (b) plot the distance in kilometers from residence to the CBD for parents and never-parents, respectively. Panels (c) and (d) graph the distance in kilometers from work to the CBD for parents and never-parents, respectively.

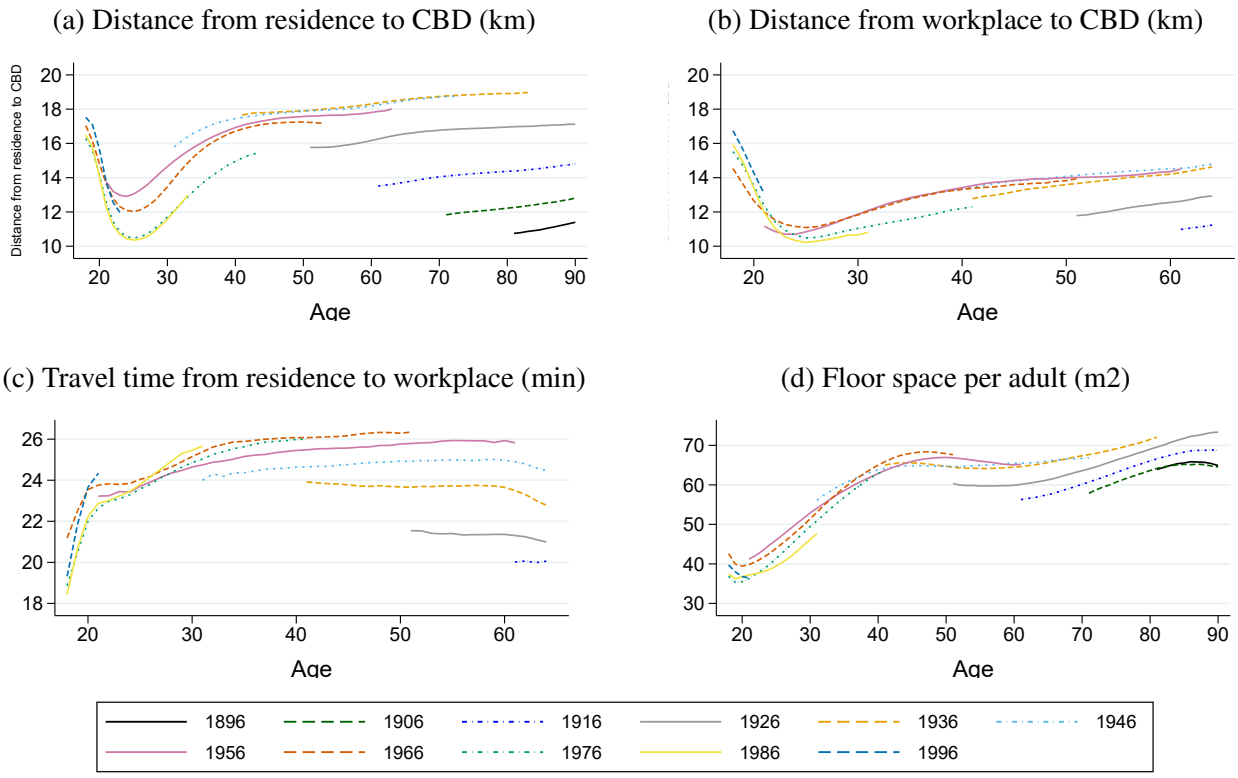
C.3 Life Events

The event-studies shown in Subsection 3.3 show the impact that certain life events may have on individuals' choices of residence and workplace location, and on their floor space consumption. We now show relevant summary statistics about the events and the sample, which are important for our estimation of the treatment effects of life events. In Table C1, we list the life events, statistics about the age at which people experience them, and the share of individuals in the estimation sample that have been treated with each event.

When we jointly estimate the effects of early life events, the sample only consider individuals that we observe at least one year as untreated by any early life event. We only consider events that are experience by at least 4% of the estimation sample, and most events are quite frequent.

The first event can either be first child or first cohabitation and the great majority experiences first cohabitation first. Out of the estimation sample, only 13.15% has a first child before ever cohabiting, and 12.6% start cohabiting and have a first child in the same year. Figure C8 shows an histogram of the distance in years between having a first child and having a first cohabitation.

Figure C6: Residence and workplace location choices over age, by cohort



Notes: This figure shows residence and location choices over the life cycle by cohort. Panel (a) plots the kilometer-measured distance from residence to the CBD against age. Panel (b) plots the kilometer-measured distance from workplace to the CBD against age. Panel (c) graphs commuting time in minutes against age. Panel (d) graphs floor space per adult in square meters against age.

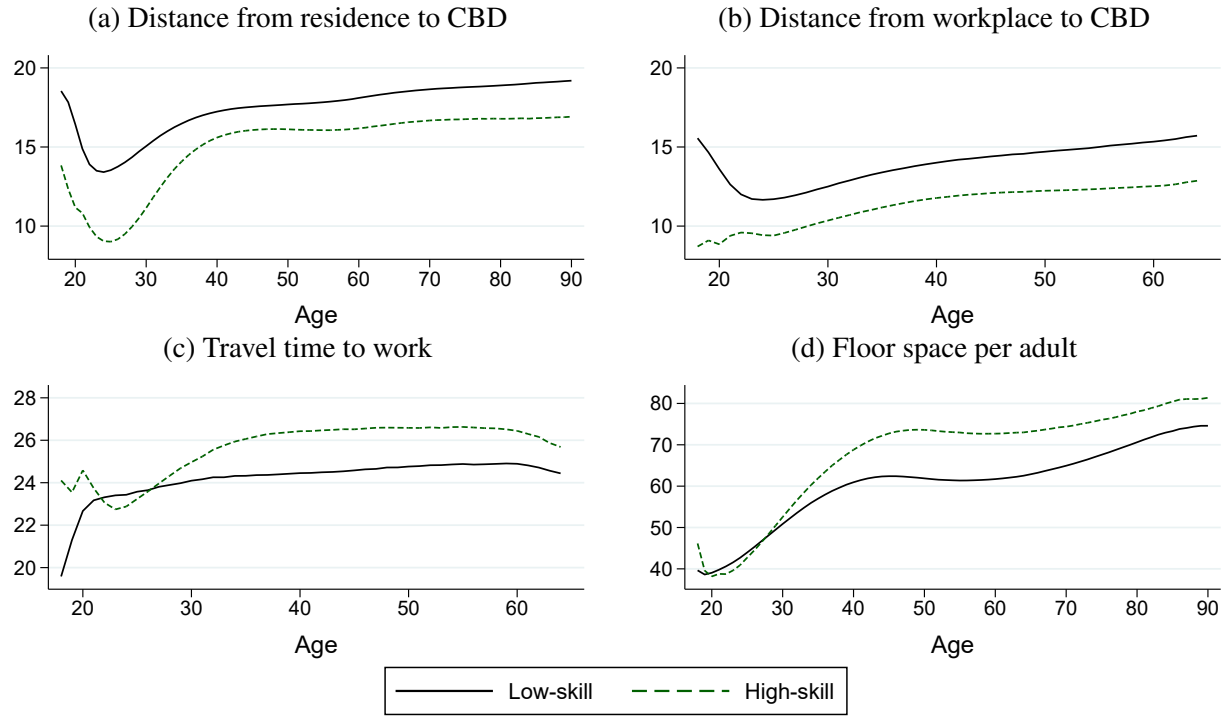
It shows that it is much more common to start cohabiting before having children, but that it is common to have a first child within the first years of the first cohabitation.

Figures C9 and C10 show the variation that exists for the age at which people experience certain events in our sample. The variation exists both in the intensive margin, relative to the years between events, and in the extensive margin, since different individuals will have experienced a different set of events throughout their lives.

C.3.1 Comparison between our strategy and OLS

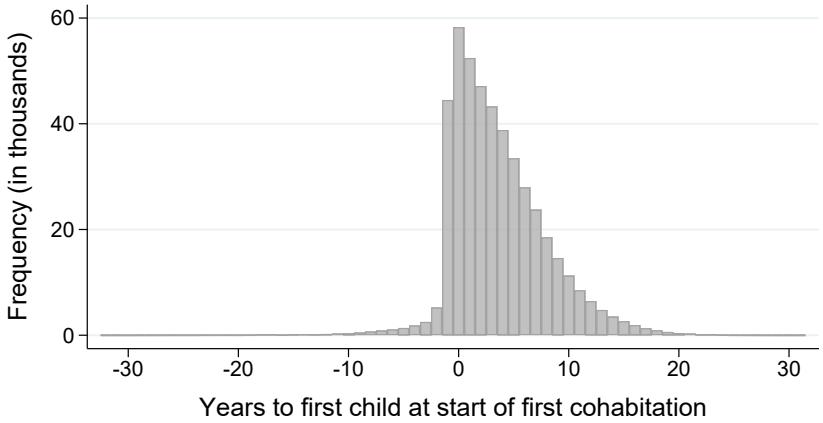
We now provide a comparison of different estimation methods and explain why the imputation method with the joint estimation of events is our preferred methodology. As explained in Appendix Section B, our

Figure C7: Residence and workplace location choices over age by skill



Notes: This figure shows residence and location choices over the life cycle by skill level. Panel (a) plots the kilometer-measured distance from residence to the CBD against age. Panel (b) plots the kilometer-measured distance from workplace to the CBD against age. Panel (c) graphs commuting time in minutes against age. Panel (d) graphs floor space per adult in square meters against age.

Figure C8: Distance between Year of First Child and Year of First Cohabitation



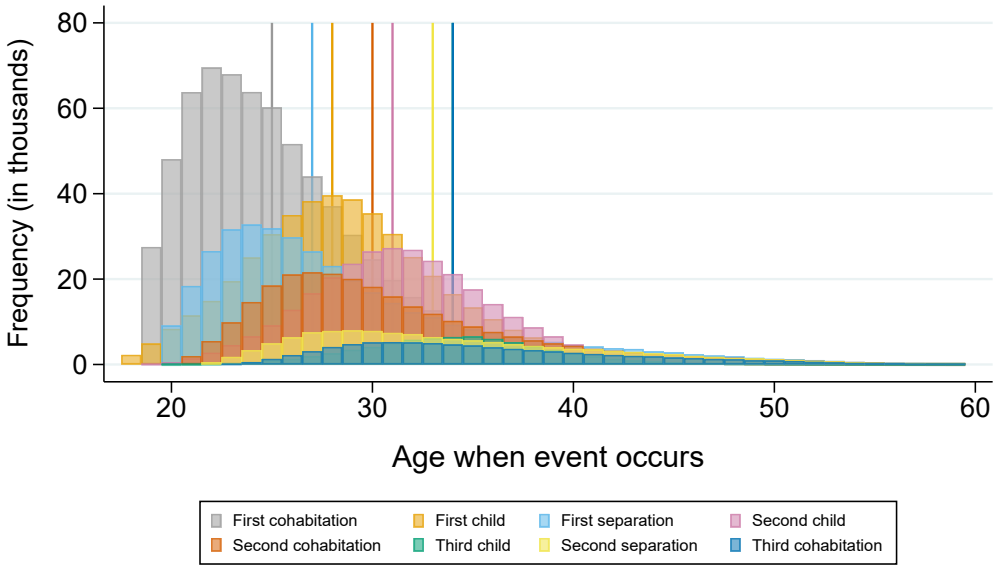
Notes: This figure shows the distribution of time elapsing between the first cohabitation and the birth of the first child.

Table C1: Age Distribution of Life Events and Sample Characteristics

Event	Age at Treatment p10 (1)	p50 (2)	p90 (3)	Number of Ever Treated Individuals (4)	Share of Ever Treated Individuals (5)
Early Life Events					
First Child	23	29	36	660,503	22.31
Second Child	25	31	38	517,545	17.48
Third Child	28	34	41	172,159	5.81
First Cohabitation	21	26	41	870,719	29.41
Second Cohabitation	25	32	51	498,638	16.84
Third Cohabitation	28	37	54	145,563	4.92
First Separation	22	31	55	804,221	27.16
Second Separation	26	36	54	241,615	8.16
Late Life Events					
Empty Nest	42	52	62	630,665	21.30
Pension	49	62	67	671,887	22.69
First Widowhood	52	70	84	201,439	6.80

Notes: The table presents the age distribution of early and life events over the life cycle. Columns (1) to (3) show the age at treatment in the 10th, 50th and 90th percentile. Column (4) shows the total number of ever treated individuals with each event and Column (5) shows the share of ever treated individuals relative to the estimation sample size, which is of 1,323,393 individuals for the early life events estimation and 2,408,353 for the late life events estimation.

Figure C9: Frequency of Early Life Events by Age

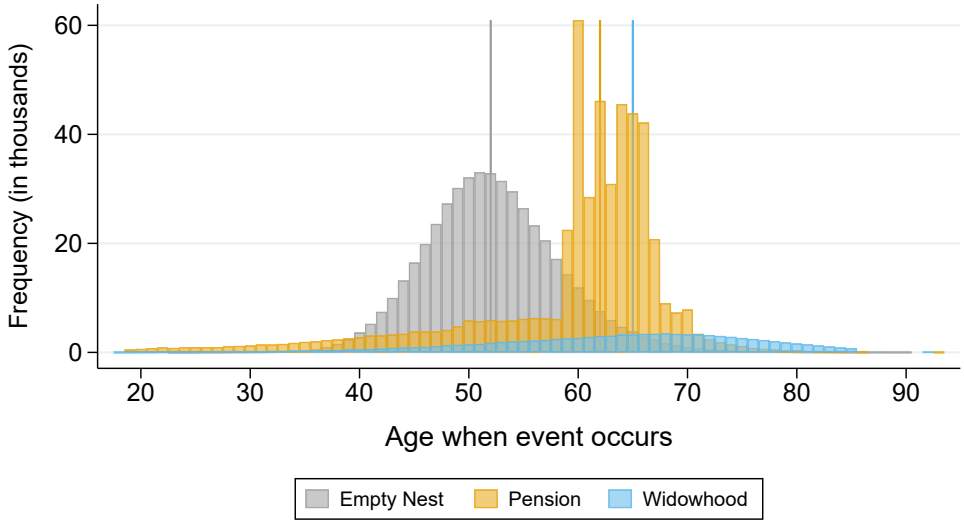


Notes: This figure plots the distribution of early life events against age.

C.3.2 Heterogeneities

We now consider the effect of two of our four major life events – first cohabitation and the birth of the first child – disaggregated by gender and skill level in turn.

Figure C10: Frequency of Late Life Events by Age



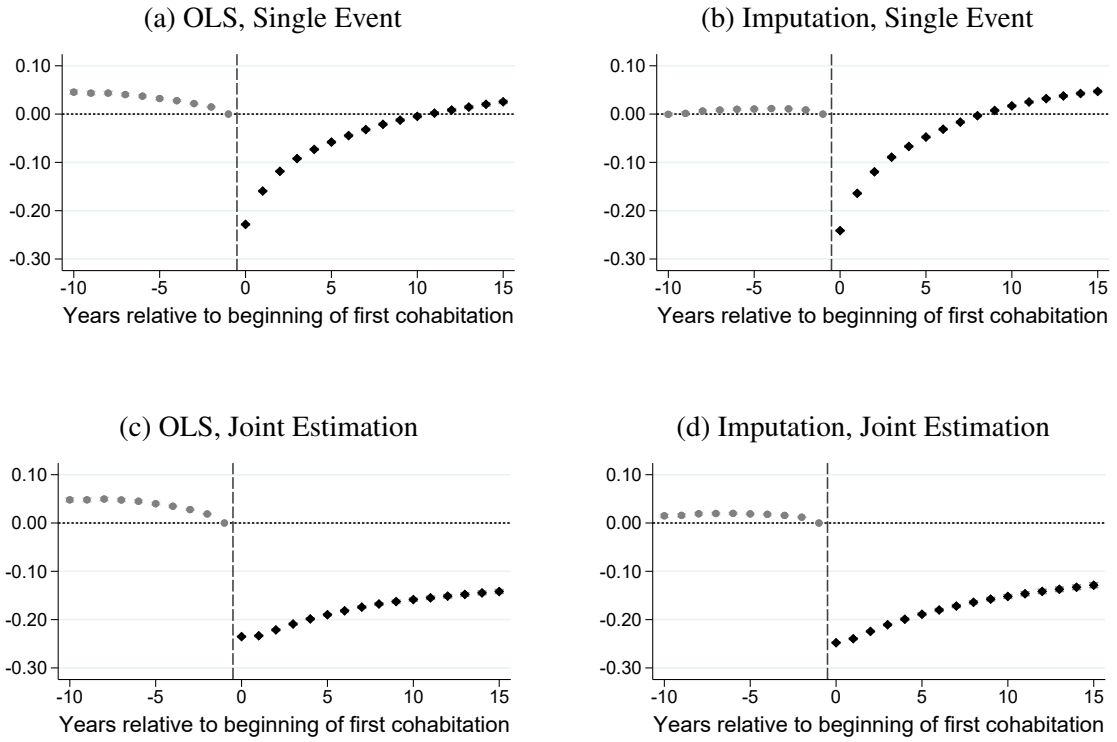
Notes: This figure plots the distribution of late life events against age.

Gender Figure C13 plots the effects of the first cohabitation disaggregated by gender. Each panel plots event-time coefficients relative to the beginning of cohabitation at year 0. Panel (a) shows the change in distance from residence to the CBD, and Panel (b) shows the change in distance from workplace to CBD, both in kilometers. Panel (c) reports the changes in commute time from residence to workplace in minutes. Panel (d) shows the change in floor space per adult in square meters. While the two bottom panels display only a slight difference between males and females, the two top panels show that with the beginning of cohabitation, females increase their distance from residence and work to the CBD relatively more than males. Although this divergence closes for distance to residence ten years after cohabitation, it further widens for distance to work.

Figure C13 plots the effects of the first cohabitation disaggregated by gender. Each subfigure plots event-time coefficients relative to the beginning of the first cohabitation at year 0. Figure C13(a) shows the change in distance from residence to the CBD, and Figure C13(b) shows the change in distance from workplace to the CBD, both in kilometers. Figure C13(c) reports the changes in commute time from residence to workplace in minutes. Figure C13(d) shows the change in floor space per adult in square meters. While the two bottom panels display only slight differences between males and females, the two top panels show that with the beginning of cohabitation, females increase their distance from residence and work to the CBD relatively more than males. Although this divergence closes for distance to residence ten years after cohabitation, it further widens for distance to work.

Figure C14 plots the effects of the first child disaggregated by gender. Each panel of this figure plots the event-time coefficients relative to the arrival of the first child at year 0. Panel (a) shows

Figure C11: Effects of First Cohabitation on Floor Space Consumption: Imputation vs. OLS

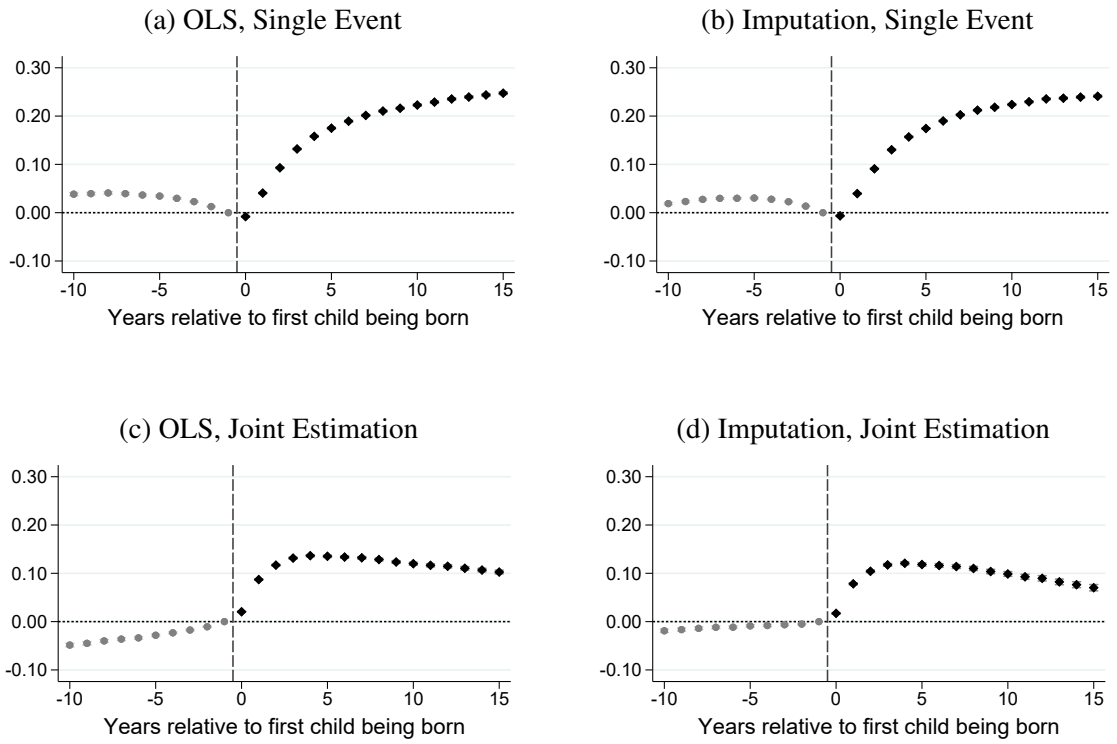


Notes: Each panel of this figure plots event-time coefficients relative to the beginning of the first cohabitation at year 0, comparing OLS to imputation estimation. Panels (a) and (b) show the change in floor space consumption relative to the years from the beginning of the first cohabitation considered as a single event and estimated by OLS and imputation, respectively. Panels (c) and (d) show the change in floor space consumption relative to the years from the beginning of the first cohabitation, considered jointly estimated by OLS and imputation respectively.

the change in distance from residence to the CBD, and Panel (b) shows the change in distance from workplace to CBD in kilometers. Panel (c) reports the changes in commute time from residence to workplace in minutes. Panel (d) shows the change in floor space per adult in square meters. For each of the four panels, the change is relatively more extreme for females than for males. Distance from residence and work to the CBD increase more for females than males after the birth of the first child. Travel time from residence to work decreases substantially more, and the gap widens in the fifteen years following the event. Floor space per adult increases more at first and maintains a constant difference in time. The combination of these four panels suggests that the arrival of the first child causes women to decentralize both their place of residence and their workplace relatively more than males, which decreases their commuting time substantially.

Skill Figure C15 plots the effects of the first cohabitation disaggregated by skill level. Each panel of this figure plots event-time coefficients relative to the beginning of the first cohabitation at year 0. Panel (a) shows the change in distance from residence to the CBD, and Panel (b) shows

Figure C12: Effects of First Child on Floor space Consumption: Imputation vs. OLS

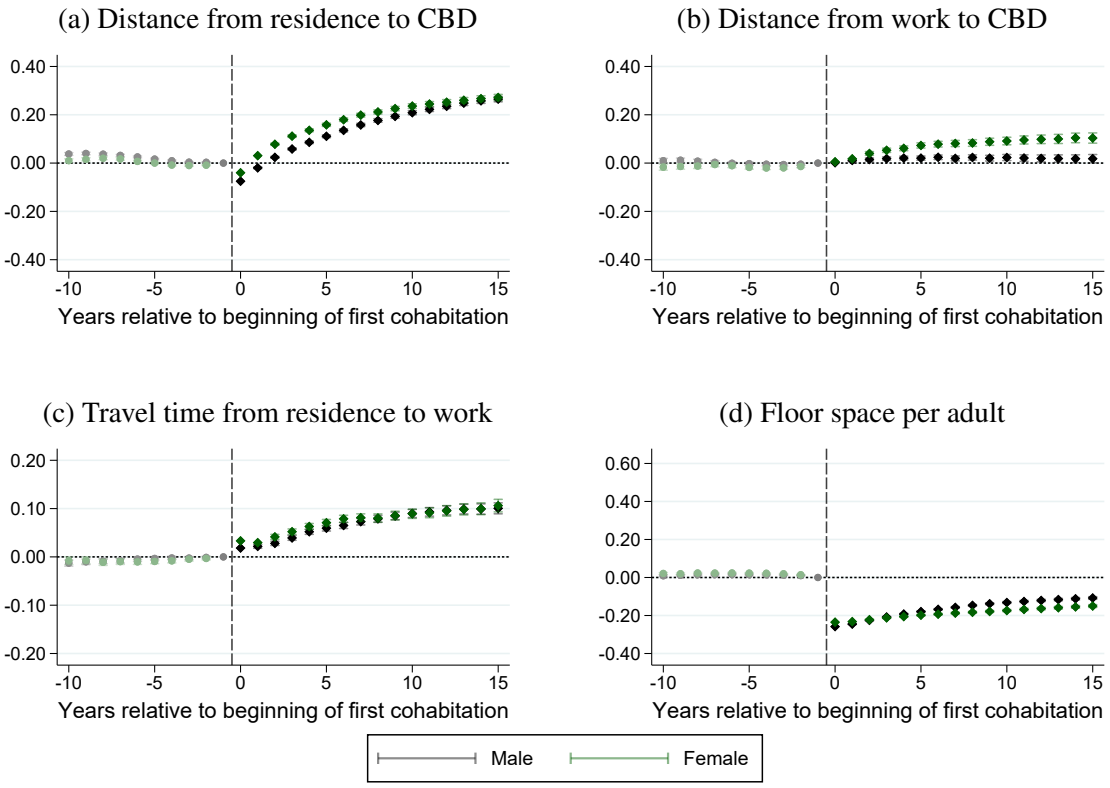


Notes: Each panel of this figure plots event-time coefficients relative to the birth of the first child at year 0, comparing OLS to imputation estimation. Panels (a) and (b) show the change in floor space consumption relative to the years from the beginning of the first cohabitation considered as a single event and estimated by OLS and imputation respectively. Panels (c) and (d) show the change in floor space consumption relative to the years from the beginning of the first cohabitation, considered jointly estimated by OLS and imputation respectively.

the change in distance from workplace to CBD in kilometers. Panel (c) reports the changes in commute time from residence to workplace in minutes. Panel (d) shows the change in floor space per adult in square meters.

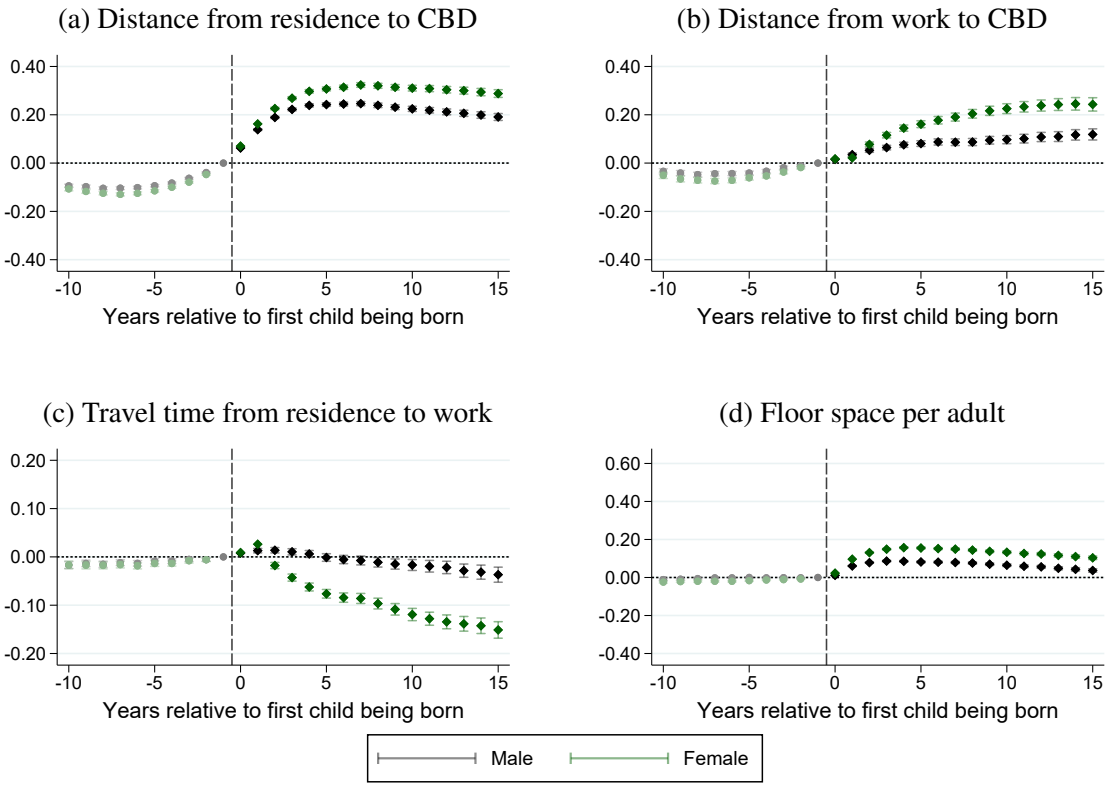
Figure C16 plots the effects of the first child disaggregated by skill level. Each panel of this figure plots the event-time coefficients relative to the arrival of the first child at year 0. Panel (a) shows the change in distance from residence to the CBD, and Panel (b) shows the change in distance from workplace to CBD in kilometers. Panel (c) reports the changes in commute time from residence to workplace in minutes. Panel (d) shows the change in floor space per adult in square meters.

Figure C13: Effects of First Cohabitation: Gender



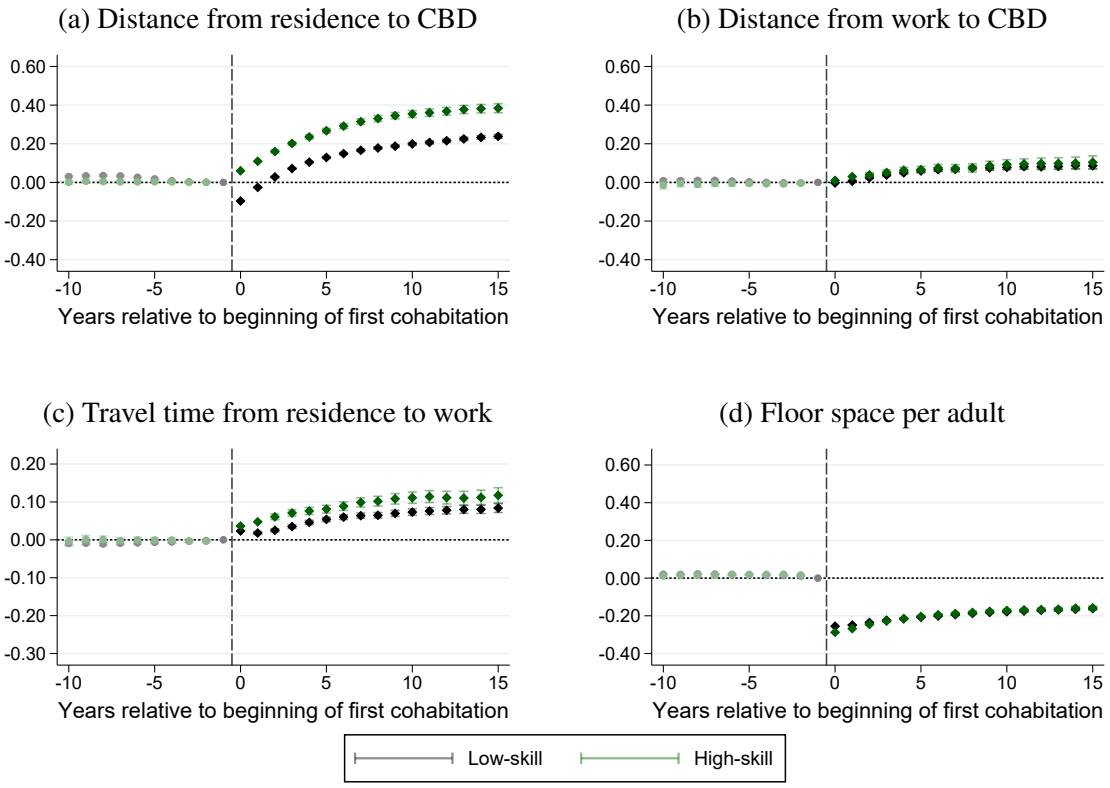
Notes: Each panel of this figure plots event-time coefficients relative to the beginning of the first cohabitation at year 0, disaggregated by gender. Panel (a) shows the change in distance from residence to the CBD, and Panel (b) shows the change in distance from workplace to CBD, both in kilometers. Panel (c) reports the changes in commute time from residence to workplace in minutes. Panel (d) shows the change in floor space per adult in square meters.

Figure C14: Effects of First Child: Gender



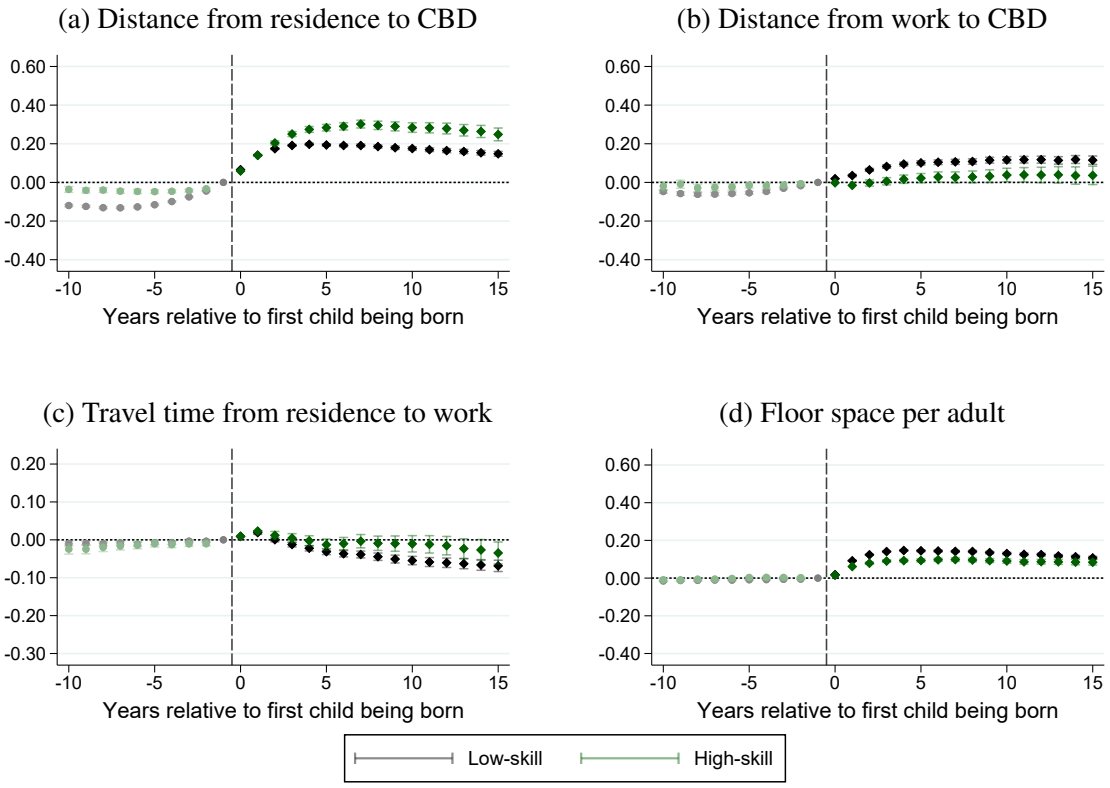
Notes: Each panel of this figure plots event-time coefficients relative to the birth of the first child at year 0, disaggregated by gender. Panel (a) shows the change in distance from residence to the CBD, and Panel (b) shows the change in distance from workplace to CBD, both in kilometers. Panel (c) reports the changes in commute time from residence to workplace in minutes. Panel (d) shows the change in floor space per adult in square meters.

Figure C15: Effects of First Cohabitation: Skill



Notes: Each panel of this figure plots event-time coefficients relative to the beginning of the first cohabitation at year 0, disaggregated by skill. Panel (a) shows the change in distance from residence to the CBD, and Panel (b) shows the change in distance from workplace to CBD, both in kilometers. Panel (c) reports the changes in commute time from residence to workplace in minutes. Panel (d) shows the change in floor space per adult in square meters.

Figure C16: Effects of First Child: Skill



Notes: Each panel of this figure plots event-time coefficients relative to the birth of the first child at year 0, disaggregated by skill. Panel (a) shows the change in distance from residence to the CBD, and Panel (b) shows the change in distance from workplace to CBD, both in kilometers. Panel (c) reports the changes in commute time from residence to workplace in minutes. Panel (d) shows the change in floor space per adult in square meters.

Table D1: Overview of Model Groups in the Extended Model

Age	Work Status	Skill	Family type
Young	W, NW	LS, HS	Single, Cohabiting, Cohabiting with Children
Senior	W, NW	LS, HS	Single, Cohabiting, Cohabiting with Children
Pensioner	NW	LS, HS	Single, Cohabiting

Notes: This table presents the combination of demographic and economic characteristics used to define groups in the extended model.

D Model

In this Appendix Section, we propose a model extension where we include the non-working adult population, as well as the workers. They do not have a workplace, commuting, and consume floor space.

D.1 Model Extension: Including the Non-Workers

We can extend our model to include non-working population and individuals with more than 65 years old, who only choose a residence location and do not have a workplace. In our sample, this will include all the individuals that do not have a workplace attached to them, the great majority being pensioners, students over age 18, and unemployed individuals. We allow our non-working model population to vary by skill, age, and family type, and this results in a total of 28 model groups, with 12 worker groups and 16 non-worker groups. We do not consider pensioners living with children, due to their infrequency in the data. Table D1 details the characteristics that the groups may have.

Preferences of Non-workers We assume that a non-worker in group r receives an exogenous wage $\bar{w}^r$, equivalent to a pension or stipend that does not vary across residence locations. It follows that the indirect utility for non-worker ρ in group r, residing in location n, depends on their exogenous wage $\bar{w}^r$, the price of residential floor space Q_n, group specific location amenities B_n^r, and an idiosyncratic amenity draw $z_n^r(\rho)$ for the residential location:

$$U_n^r(\rho) = \frac{B_n^r \bar{w}^r z_n^r(\rho)}{(P_n)^{\alpha^r} (Q_n)^{1-\alpha^r}} \quad 0 < \alpha^r < 1 \tag{D.1}$$

The idiosyncratic taste shocks $z_n(\rho)$ influence the residential location choice, and are independent draws from a Fréchet distribution with cumulative distribution function. $G(z) = e^{-(z)^{-\varepsilon^r}}$. The Fréchet shape parameter $\varepsilon^r > 1$ controls the dispersion of utility and varies by non-worker group r. Non-workers choose their residence location as to maximize their utility. The properties of the

Fréchet distribution allow us to write the probability that a non-worker from group r chooses to live in n as:

$$\lambda_n^r = \frac{L_{Rn}^r}{L^r} = \frac{(B_n^r)^{\varepsilon^r} \left(Q_n^{1-\alpha^r}\right)^{-\varepsilon^r}}{\sum_{k\in\mathbb{N}} (B_k^r)^{\varepsilon^r} \left(Q_k^{1-\alpha^r}\right)^{-\varepsilon^r}} \quad n \in \mathbb{N} \tag{D.2}$$

where L_{Rn}^r is the measure of non-workers from group r residing in location n, and L^r is the number of individuals in group r.

Floor space market clearing Lastly, the market clearing condition for residential floor space expands to include non-working agents in each location i:

$$H_{Ri} = \sum_{o\in\mathbb{O}} \sum_{f\in\mathbb{F}} (1-\alpha^{of}) \frac{v_i^{of} L_{Ri}^{of}}{Q_i} + \sum_{r\in\mathbb{R}} (1-\alpha^r) \frac{\bar{w}^r L_{Ri}^r}{Q_i} \tag{D.3}$$

General Equilibrium The equilibirum conditions for the extended model remain similar to the original one. Given the extended model's structural parameters $\{\phi^{of}, \alpha^{of}, \alpha^r, \beta^H, \beta^o, \varepsilon^{of}, \varepsilon^r\}$, the city's populations by group $\{L^{of}, L^r\}$, the income of the non-workers per group $\{\bar{w}^r\}$, and the exogenous vector of location specific characteristics $\{A_i, B_{ni}^{of}, B_n^r, H_{Fi}, H_{Ri}\}$, the general equilibrium of the model is referenced by the vector of endogenous variables $\{L_{Ri}^{of}, L_{Ri}^r, w_i^o, L_{Fi}^{of}, Q_i, q_i\}$ determined by the following equations: Worker residential choice probabilities (8); non-worker residential choice probabilities (D.2); profit maximization and zero profit condition (10); commuter market clearing for each worker type (11); residential floor space clearing (14); commercial floor space clearing (15).

E Model Estimation

E.1 Calibrating the Group Sizes

Before describing the details of the model estimation, it is important to explain how the groups we consider in the model appear in the data. Our model population is Copenhagen's working-age population, defined as individuals aged 18-65. We classify individuals along three dimensions: life stage, skill level, and family type. The groups are described in Appendix Table E1. We use the terms life stage and age group interchangeably since they are closely linked. Workers can be young (age 18 to 45) or senior (45 to 65). Those older than Denmark's state pension age (65) are excluded from the modeled population and treated as pensioners. Appendix Figure E1 shows each age group's share of the CMA adult population and how stable these shares are over time.

The second relevant dimension to classify individuals is skill. Individuals are divided into low- and high-skill, where high-skilled is used to describe someone with at least some tertiary education.[40] Students are not considered in this skill classification and excluded from the model population. As shown in Appendix Figure E2, the share of high-skilled workers in the model population rose from 22.5% in 1986 to 49.1% in 2017.

Lastly, we classify individuals by family type: single, cohabiting with a partner, and cohabiting with a partner and children. We exclude single parents from the model since they constitute a small share of the population as shown in Appendix Figure E3. As with the age groups, the distribution of family types is very stable over time.

Table E1: Relevant dimensions for defining the model groups

Life stage	Skill	Family type
Young worker	LS, HS	Single, Cohabiting, Cohabiting with Children
Senior worker	LS, HS	Single, Cohabiting, Cohabiting with Children

Notes: This table details the dimensions we consider for defining the groups in the model. Individuals are in different life stages, skill group and family type.

Our model considers 12 different groups. These groups account for 41.74% of the CMA adult population during the period considered, and for 41.69% of the CMA population in 2011, the year we use to estimate the model.

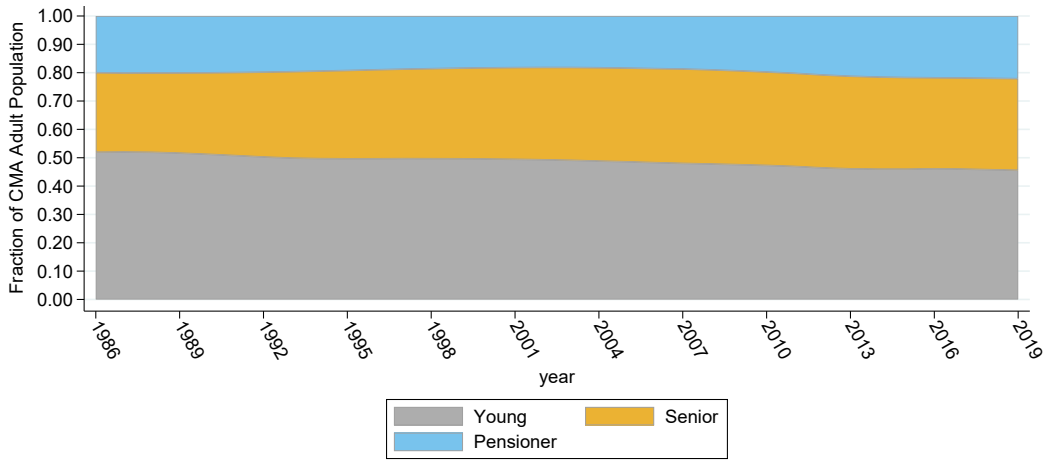
Figure E1 plots the evolution of age group shares in the CMA population from 1986 to 2019. The distribution remains relatively stable over time, with a slight decline in the young population and a modest increase in seniors.

[40] In the first years of the panel the level of skill is sometimes missing. To include all individuals and still respect the model heterogeneity, we predict missing skill using the income percentile.

Figure E2 dislays the evolution of labor group shares in the CMA adult population from 1986 to 2019. The figure suggests that the share of high-skilled individuals has increased, particularly among the young, while the share of low-skilled young workers has declined steadily.

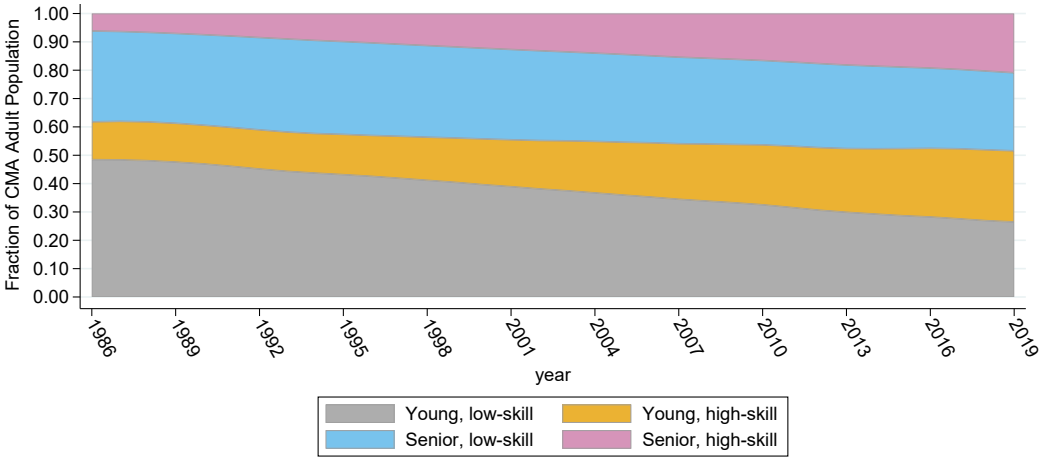
Figure E3 shows the evolution of family group shares in the CMA adult population from 1986 to 2019. Over time, the share of high-skilled individuals has increased, particularly among the young, while the share of low-skilled young workers has declined steadily. The distribution remained broadly stable over time, with a slight increase in the share of single adults without kids.

Figure E1: Evolution of age group shares over time



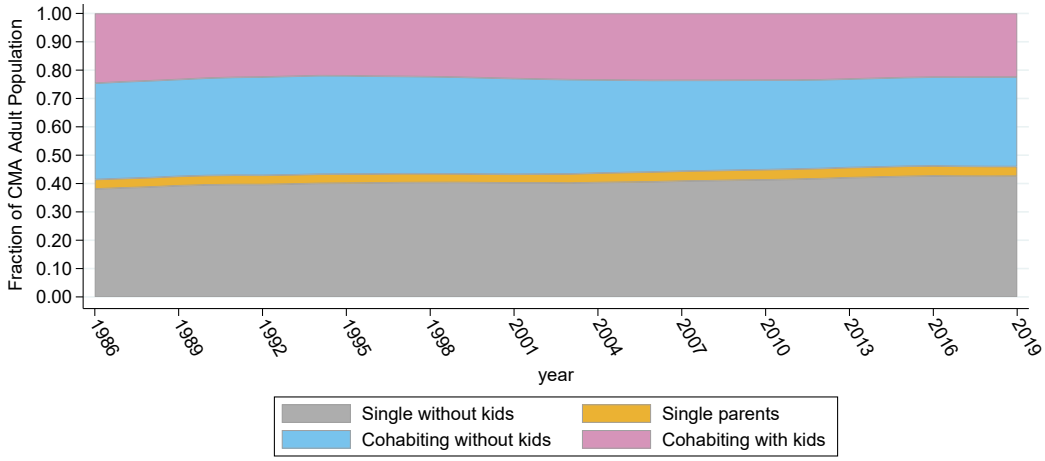
Notes: The figure plots the evolution of age group shares in the CMA adult population from 1986 to 2019.

Figure E2: Evolution of labor group shares over time



Notes: The figure shows the evolution of labor group shares in the CMA adult population from 1986 to 2019.

Figure E3: Evolution of family group shares over time



Notes: The figure shows the evolution of family group shares in the CMA adult population from 1986 to 2019.

E.2 Parameter Estimation

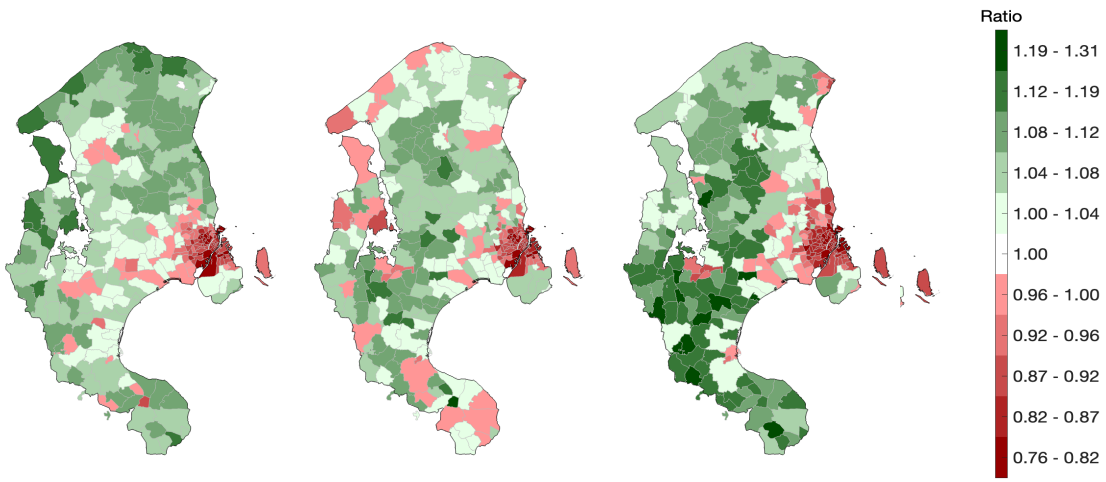
Housing Expenditure Shares For the model calibration, we estimate group-specific housing expenditure shares, α_{of}. In our population panel, we observe aggregate household income net of taxation as well as residential location and floor space consumption. Using the universe of property transactions, we establish a local index of residential floor space prices in each year using the method of Ahlfeldt et al. (2021). To ensure that we are capturing location choices based on current market prices, we then restrict the sample to those that moved residence in the previous year. Lastly, we calibrate an annual yield to target an aggregate housing expenditure share of 30% for the total model population [41], and then estimate average housing expenditure shares for each group. The results are presented in Table 1.

Residential Amenities Appendix Figure E4 compares the estimated residential amenities along three different demographic dimensions. For each panel, we use the sub-group average weighted by the residential population in each parish to calculate the residential amenity for each aggregation and then take the ratio of the two groups. Panel (a) shows the residential amenities of the senior workers relative to the young workers, and indicates that the young have a higher preference for central locations. Panel (b) compares parents with non-parents, and indicates that parents prefer locations outside of the city center. Panel (c) compares couples with singles, and indicates that singles prefer to locate in central locations.

[41] See the 2021 Productivity report in https://dors.dk/vismandsrapporter/produktivitet-2021.

Figure E4: Residential Amenity Comparisons

(a) Senior vs. Young (b) Parents vs. Non-Parents (c) Couples vs. Singles



Notes: Appendix Figure E4 maps the relative amenities between groups of different characteristics, and we observe that young, single, and non-parent groups have higher amenity values in the CBD.

F Demographic Counterfactuals

In this section, we describe how the demographic data for the Greater Tokyo Area[42] in 1980 and 2020 (presented in Table 3) was constructed. The main data sources for this table were the 1980 and 2020 Japanese censuses.[43]

We then detail how we simulate the observed demographic changes in the Japanese data using our data in Copenhagen. We analyze demographic changes along three different dimensions: population aging, declining fertility, and delayed timing of marriage resulting in more singles in the population. Lastly, we simulate the effect of all three of these changes together.

F.1 Japanese Census Data

Age Composition Table 2 of the 1980 Japanese census provides data on the total population by age (in single years) for each Japanese prefecture[44]. Using this data, we collect the number of people in the 18-44, 45-64 and 65+ age brackets and construct shares for these categories. Note that our dataset excludes everyone under 18 and people whose age is listed as unknown in the census[45].

Table 2-1-1 of the 2020 Japanese census provides the same data for all Japanese prefectures in 2020[46]. We use the same process to construct shares of the 18-44, 45-64 and 65+ age brackets in the total adult (18+) population in the Greater Tokyo Area.[47]

Coupling Status The Japanese censuses does not provide data on unmarried couples in the Japanese populace. As such, we use marital status data to determine the number of singles and couples in Japan – all individuals with marital status as "Married" are classified as in a couple,

[42]There are several definitions of what constitutes the Greater Toyko Area. We adopt the definition used by the Japanese Government in their National Capital Region Development Plans (NCRDPs), which define the Greater Tokyo Area as all the prefectures in the Kantō region (Chiba, Gunma, Ibaraki, Kanagawa, Saitama, Tokyo Metropolis and Tochigi) and the prefecture of Yamanashi of the neighboring Chūbu region (Jain and Arai 2019).

[43]Note that data for each demographic dimension has been collected from a different census dataset. For each demographic dimension, individuals whose demographic information is unknown have been excluded from the analysis. This leads to different figures for the total adult (18+) population of the Greater Tokyo Area depending on the dataset used. For instance, the marital status dataset suggests there are 34,825,562 adults in the Greater Tokyo Area in 2020, while the age dataset suggests there are 37,102,147. This is because there are more individuals whose marital status is unknown than individuals with unknown ages in the Greater Tokyo Area. Shares for each demographic dimension have been constructed using the total 18+ population figure implied by that dataset only.

[44]This dataset can be found here: Table 2 of the 1980 Japanese Census.

[45]There are 35,905 people whose age is listed as unknown in the Greater Tokyo Area. This represents less than 0.2% of the known adult population in the Greater Tokyo Area in 1980.

[46]This dataset can be found here: Table 2-1-1 of the 2020 Japanese Census.

[47]There are 1,144,789 people whose age is listed as unknown in the Greater Tokyo Area. This represents around 3% of the known adult population in the Greater Tokyo Area in 2020.

while all other individuals are classified as single (excluding people whose marital status is listed as "Unknown").[48]

Table 5 of 1980 Japanese census[49] provides marital status data for all 15+ individuals in each Japanese prefecture by 5 year age groups (15-19, 19-24, etc.), while Table 4[50] provides marital status data for all 15+ individuals across Japan by age (in single years). Using the first dataset, we collect marital status data for the 20-44, 45-64 and 65+ age groups in the Greater Tokyo Area. To include the 18-19 age group, we use the Japan-wide data to identify the fraction of singles and couples in the 15-19 age group that are 18 or 19 years old.[51]We multiply these fractions by the number of singles and couples in the 15-19 category in the Greater Tokyo Area to estimate the number of singles and couples aged 18-19 in the Greater Tokyo Area. We then aggregate across the age groups to obtain the total number (and share) of 18+ singles and couples in the Greater Tokyo Area in 1980.[52]

Table 4-1 of the 2020 Japanese census provides marital status data for all 15+ individuals in each Japanese prefecture by age (in single years)[53]. We use this dataset to identify 18+ singles and couples in the Greater Tokyo Area in 2020.[54]

Parental Status The Japanese censuses do not provide data on the number of parents in Japan. Instead, they classify the population into different households depending on their living structure at home. We use this household classification, along with some assumptions, to estimate the number of parents and non-parents in Japan in 1980 and 2020. Note that our definition of parent refers to an individual who lives with their biological child(ren) who is (are) under 18 years of age.

The 1980 census broadly classifies households into 2 types - "ordinary households" and "quasi-households". An "ordinary household" is defined as a group of persons sharing living quarters (typically families) and living expenses as well as a person who lives by himself occupying a dwelling house, while all other households are classified as "quasi-households"[55].

There are 16 types of "ordinary households" in Japan, which are classified based on the relationship of household members to the household head. Based on the household description, we

[48]Individuals are classified as Never Married, Married, Widowed, Divorced or Unknown.

[49]This dataset can be found here: Table 5 of the 1980 Japanese Census.

[50]This dataset can be found here: Table 4 of the 1980 Japanese Census.

[51]38% of all singles and 85% of all married couples in the 15-19 category are 18 or 19 years old in the Japan-wide data.

[52]Note that less than 0.5% of the 15+ population in the Greater Tokyo Area have their marital status listed as Unknown in 1980. These individuals have been excluded from the data.

[53]This dataset can be found here: Table 4-1 of the 2020 Japanese Census.

[54]Note that around 3% of the 18+ population in the Greater Tokyo Area have their marital status listed as Unknown in 2020. These individuals have been excluded from the data.

[55]Quasi-households include students in school dormitories, inpatients of hospitals, inmates of social institutions, persons in camps of Self Defence Forces, inmates of reformatory institutions, single persons in company's dormitories and others.

classify household types as those containing parents (e.g: "Married couple with children") and those that do not (e.g: "Married couple only"). For the household types that contain parents, we assume the number of parents per household based on the household description. For instance, the "Married couple with children" category would have 2 parents per household, while the "Mother with child(ren)" category would have only 1 parent per household. The full list of "ordinary households" and their classification into parents and non-parents is provided in Table F1.

Table F1: Ordinary Household Classification

Ordinary Household Type	Contains parents	No. of parents assumed per household
Married couple only	No	-
Married couple with children	Yes	2
Father with child(ren)	Yes	1
Mother with child(ren)	Yes	1
A couple with their parents	No	-
A couple with their parent	No	-
A couple with children and parents	Yes	2
A couple with children and a parent	Yes	2
A couple with relatives other than children	No	-
A couple with their children and relatives other than parents	Yes	2
A couple with their parents and relatives other than children	No	-
A couple with children, parents and other relatives	Yes	2
Brother or sisters only	No	-
Other relatives households not elsewhere classified	No	-
Non-relative households	No	-
One-person households	No	-

Notes: This table lists household types and indicates whether they include parents, along with the number of parents assumed per household when applicable.

Table 24 of the 1980 Japanese census[56] lists households, household members and household members under 18 years for each Japanese prefecture. We first remove the number of household members under 18 from the total number of household members for each household type to keep only the 18+ population in the Greater Tokyo Area. Using this dataset and our assumed number of parents per category as presented in Table F1, we calculate the number of parents in the Greater Tokyo Area by multiplying the number of households with the number of parents assumed per household for each "ordinary" household type. We then subtract the number of parents from the total number of 18+ household members for each household type to obtain the number of non-parents. Additionally, all individuals living in quasi-households are classified as non-parents, except for students living in school dormitories, who are excluded from the analysis.[57] We then aggregate over "ordinary" and "quasi-households" to obtain the total number of parents and non-

[56] This dataset can be found here: Table 24 of the 1980 Japanese Census.

[57] This is because these individuals are likely under the age of 18. The data on individuals living in quasi-households can be found here: Table 19 of the 1980 Japanese Census.

parents in the 18+ population of the Greater Tokyo Area in 1980.[58]

The 2020 census classified households into "private" and "institutional" households, which are broadly similar to the 1980 classifications of "ordinary" and "quasi-households" respectively.[59] The family structure classification of "private" households is identical to the 1980 classification of "ordinary" households, so we use the same assumptions listed in Table F1 to classify individuals as parents and non-parents in 2020.

Tables 9-1 and 9-3 of the 2020 Japanese census[60] list the number of households and household members (and household members under 18 years) respectively for each "private" household type in all Japanese prefectures. We follow the same approach as the 1980 data to filter out individuals under 18 years and classify the 18+ population into parents and non-parents in the Greater Tokyo Area. Note that, like 1980, all individuals living in institutional households are classified as non-parents, except for students living in school dormitories, who are excluded from the analysis.[61]

F.2 Counterfactuals with Tokyo's Demographic Changes

Simulating the Effect of Japan's Demographic Changes We calibrate our model to three demographic trends observed in Japan's population from 1980 to 2020, as shown in Table 3. Each pair of columns shows how the demographic composition changed between 1980 and 2020 along three dimensions.

For the population aging counterfactual, we observe a 15% point increase in the share 45-64-year-olds in Japan from 1980 to 2020. We impose this same demographic shift in the age groups of our model, holding constant the distribution of skills and family type. Specifically, within each family-skill cell, we shift the population from the young and to the senior group. For example, we increase the Senior-Single-Low-skill share by the same absolute amount that we decrease the Young-Single-Low-skill share, with each transfer proportional to the specific Seniors's initial size. That way the overall population of Single and Low-skill and their shares in the population will remain unchanged. Applying this procedure across all six senior groups delivers the targeted 15 pp increase in the senior share without altering the skill- or family-type distributions.

We apply the same method for the declining fertility counterfactual. The share of parents in Japan fell by 24 pp from 1980 to 2020. To replicate this in Copenhagen, we shift population mass in each age–skill cell from couples with children into the corresponding couples-without-children cell. After performing this transfer for all four parental groups, we achieve the 24 pp decline in the

[58] Note that we exclude people whose household type is unknown from this analysis.

[59] The key difference is that private households also include people residing in a boardinghouse and people residing in a dormitory for unmarried employees of a company.

[60] These datasets can be found here: Table 9-1 of the 2020 Japanese Census and Table 9-3 of the 2020 Japanese Census.

[61] The data on individuals living in quasi-households can be found here: Table 6-2 of the 2020 Japanese Census.

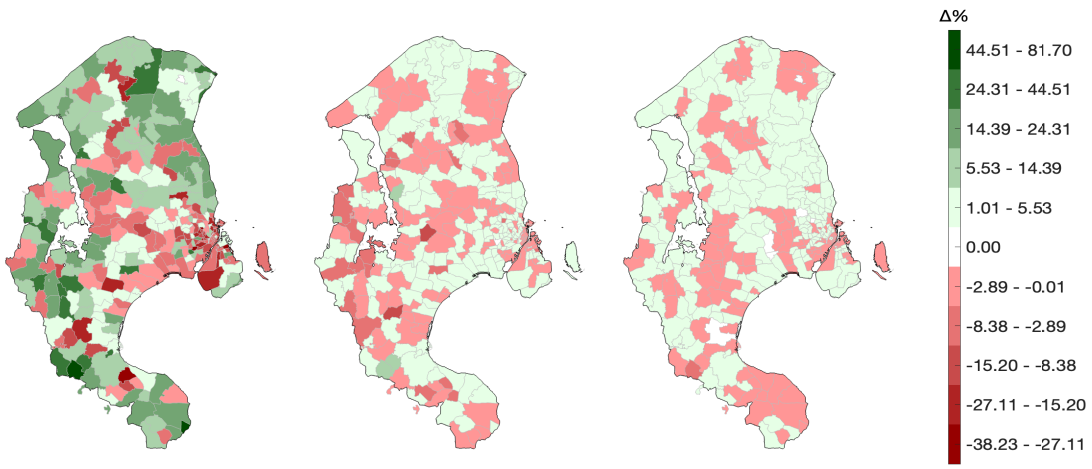
parental share while preserving the age, skill, and single-adult groups.

For the increasing singles counterfactual, we observe an 8% increase in the share of singles in Japan from 1980 to 2020. We simulate this by shifting, within each age–skill cell, mass from couples without children into the single-adult group. Performing this for each of the four single-adult cells yields the 8 pp increase in singles without affecting the age, skill, or parental margins.

Lastly, we simulate a combined counterfactual which implements all three of these demographic changes together. To implement the combined counterfactual, we sequentially impose the three demographic changes described above. Since each of the individual changes preserves the shares of the other two demographic dimensions, the order in which we implement the three changes has no effect on the resulting population distribution. [62]

Additional Demographic Counterfactual Results Figure F1 shows the effect of the three demographic changes on workplace employment. The effects on employment are not as straightforward as the effects on residential populations. Panel (a) shows that population aging predominantly decentralizes employment, while panels (b) and (c) show that the fertility and singles counterfactuals have more of a mixed effect on employment location choices.

Figure F1: Counterfactual Changes in the Workplace Employment

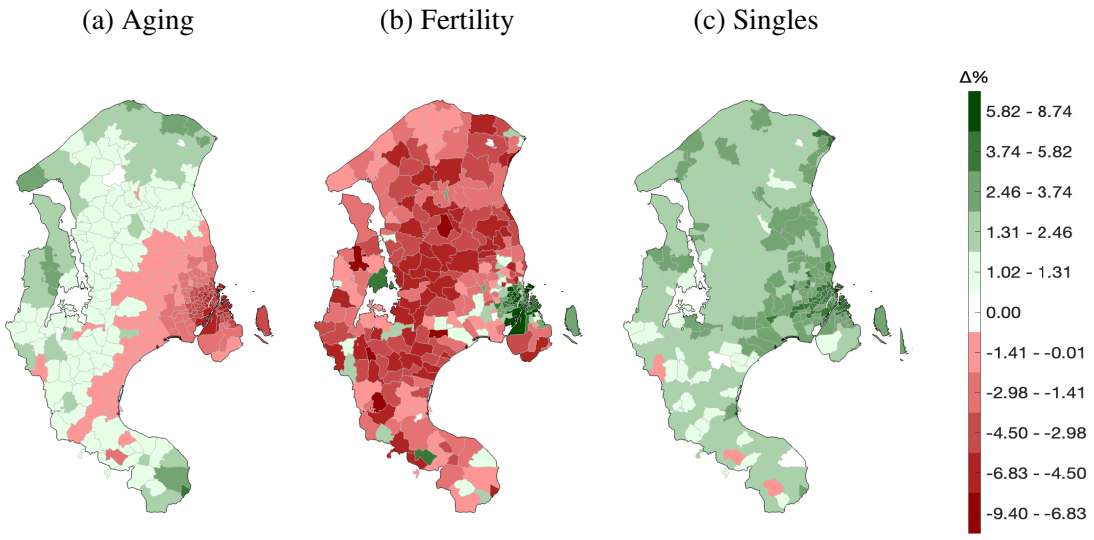


Notes: The figure presents counterfactual changes in workplace employment under three demographic scenarios: (a) aging, (b) fertility decline, and (c) increase in single households. Each panel shows the spatial variation in percentage changes in workplace employment across areas.

[62] For completeness, there are slight differences in the results when the order is changed (+/- 1 person for some groups), which are due to differences in the rounding after implementing each change.

Figure F2 shows the effect of the three demographic changes on residential floor space prices. Panel (a) shows that population aging makes the suburbs more expensive relative to the CBD, as the increase in the residential population in the suburbs drives up demand for residential floor space. Panel (b) shows that the reverse is true when fertility declines; as the residential population shifts into the center this increases the demand for residential floor space which drives up the price since the supply is fixed. Lastly, panel (c) shows that the presence of more singles in the population drives up floor space prices in nearly all locations since singles demand larger quantities of floor space than couples. Even though prices increase in all locations, it is also clear that the effect is largest in the areas which experience the most population growth.

Figure F2: Counterfactual Changes in Residential Floor space Prices

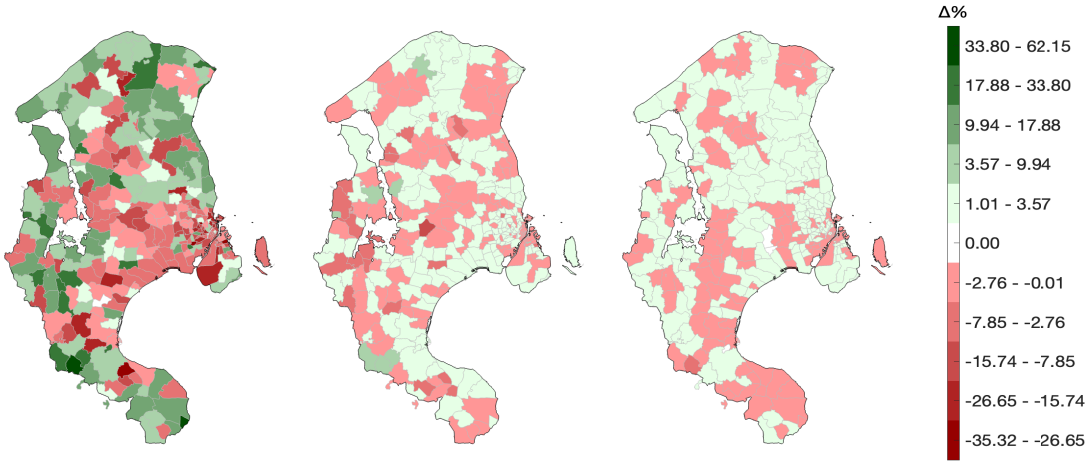


Notes: The figure presents counterfactual changes in residential floor space prices under three demographic scenarios: (a) aging, (b) fertility decline, and (c) increase in single households. Each panel shows spatial variation in percentage changes across areas.

Figure F3 shows the effects of the three demographic changes on commercial floor space prices. Similarly to the relationship between residential population and floor space prices, the changes in commercial floor space prices mirror those of workplace employment. In parishes that experience an increase in workplace employment, the price of commercial floor space increases as the firms located in this parish increase their demand for floor space proportionately to their increased workforce. Since each parish has unique production function shares for the four worker types, it is not always that case that the floor space price increases when the employment population increases.

Figure F3: Counterfactual Changes in Commercial Floor space Prices

(a) Aging (b) Fertility (c) Singles

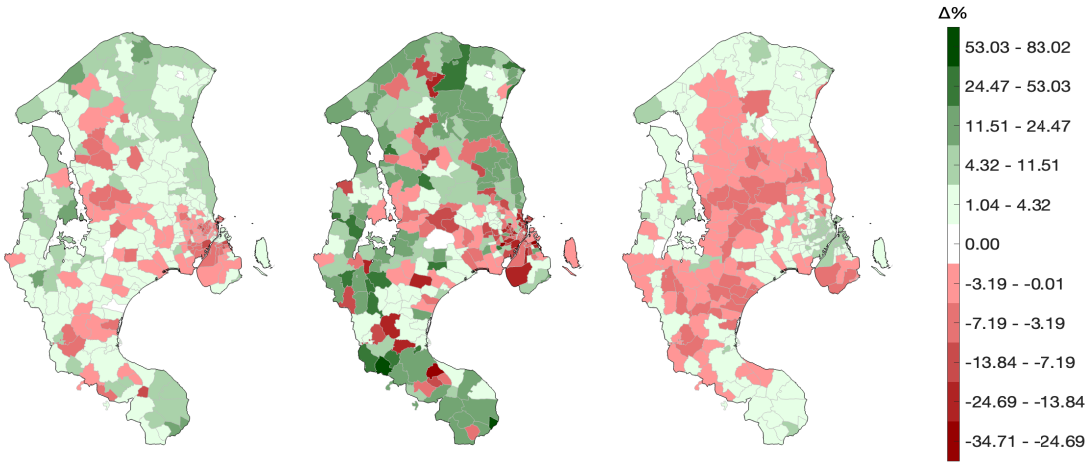


Notes: The figure presents counterfactual changes in commercial floor space prices under three demographic scenarios: (a) aging, (b) fertility decline, and (c) increase in single households. Each panel shows the spatial variation in percentage changes in commercial prices across areas.

Lastly, Figure F4 shows the combined counterfactual effects of the three demographic changes experienced in Tokyo from 1980 to 2020. Panel (a) shows that the combination of these effects predominantly decentralizes the residential population, although this effect is not as strong as in the individual counterfactuals. Panel (b) shows that workplace employment largely decentralizes in response to the demographic changes. Panel (c) shows that residential floor space prices increase in both the center and the very periphery of the city, with the median experiencing a decline in prices. This is likely due to the changing composition of the population in these different regions. Singles and couples without children prefer the CBD more than couples with children, and singles demand more floor space than couples. Both of these effects contribute to the increase in residential floor space prices in the center. Conversely, the senior population prefers the periphery to the young population, and this contributes to the increase in floor space prices in the periphery.

Figure F4: Tokyo Counterfactual Changes

(a) Residential Pop. (b) Workplace Empl. (c) Res. Floor space Price



Notes: The figure presents counterfactual changes in Tokyo under three outcomes: (a) residential population, (b) workplace employment, and (c) residential floor space prices. Each panel displays the spatial distribution of percentage changes across areas.