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

20 March 2025

Preliminary and Incomplete Draft

Abstract

Large literatures have studied the effect of age and life events, such as marriage and childbirth, on labor supply, earnings and consumption. In contrast, we know little about how age and life events affect location choices, despite vast differences in house prices and amenities across locations. This paper uses newly constructed individual level panel for Denmark covering more than 30 years to provide several new stylized facts on how workplace and residence choices, commuting, and residential floor space consumption change with age and life events. We combine the reduced form evidence with a quantitative urban model to determine what mechanisms are driving the strikingly different location choices of different groups. We use model counterfactuals to determine the relative importance of alternative mechanisms, such as differences in housing expenditure shares, commuting costs or residential amenities, for the strikingly different locations choices of different groups. Finally, we use the estimated model to evaluate the impact of population aging, changing cohabitation rates and lower fertility on the geography of cities.

JEL Codes: R21, R23, J12, J13, J14.
Keywords: Spatial sorting, quantitative spatial models, life-cycle, 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 Carolyn Diamond, Alexander Hansen, Kalle Kappner, Torge Marxsen, and Ashok Manandhar for outstanding research assistance. 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

Large literatures have studied the effect of age and life events, such as marriage and childbirth, on labor supply, earnings and consumption. In contrast, we know surprisingly little about how age and life events affect location choices. This is a major gap in our understanding as housing expenditure is the largest single expenditure item for many households and house prices, residential amenities and commuting times vary substantially across locations. One reason for this gap in our understanding is that researchers typically do not have access to data that tracks workers at a fine geographic scale over long periods of time.

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. We use the 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 these questions. 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 census data where it is often difficult to link individuals across censuses, we observe mobility at much lower temporal frequencies, and typically for much larger geographical units than 100 by 100 meter grid cells.

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 declines 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 20% 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 jointly estimate the effects of all early life events and in a separate regression jointly estimate the effects of the late life events.[2] By doing so, we account for the potential correlation between life events within the same life stage.

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 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.

[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 that the person fixed effects and age effects for the late life events can be identified from people who have been treated with an arbitrary number of early life events but no late life events yet.

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, gyms, parks, or playgrounds 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, the presence of children and who may or may not work. 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 and a further five different types of non-working population.

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 in which 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 three key parameters on sorting, which are first-order approximations to the full general equilibrium effect. 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 rules out 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. In work in progress, we correlate the model estimated amenity value that different groups attache to a location to proxies for amenities, such as the number of restaurants and bars or the amount of green space and playgrounds.

In a final step, we use the quantitative model to assess the impact of key demographic changes on the geography of cities. In particular, we estimate model counterfactuals for changes in fertility (i.e. the number of people with children), the share of single households and population aging. 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 both fewer children and more single households increase demand for space in the center of cities, this effect is offset by population aging, which reduces demand for central areas.

Related Literature Our paper is related to a number of literatures. First, our paper is related to the large literature that considers the effect of age on economic behavior. 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

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

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.

Second, a large literature 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.

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). We build on the strand 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 types of both workers and non-working population rather than the small number of worker types considered in the existing literature.

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.[4] 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

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

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.

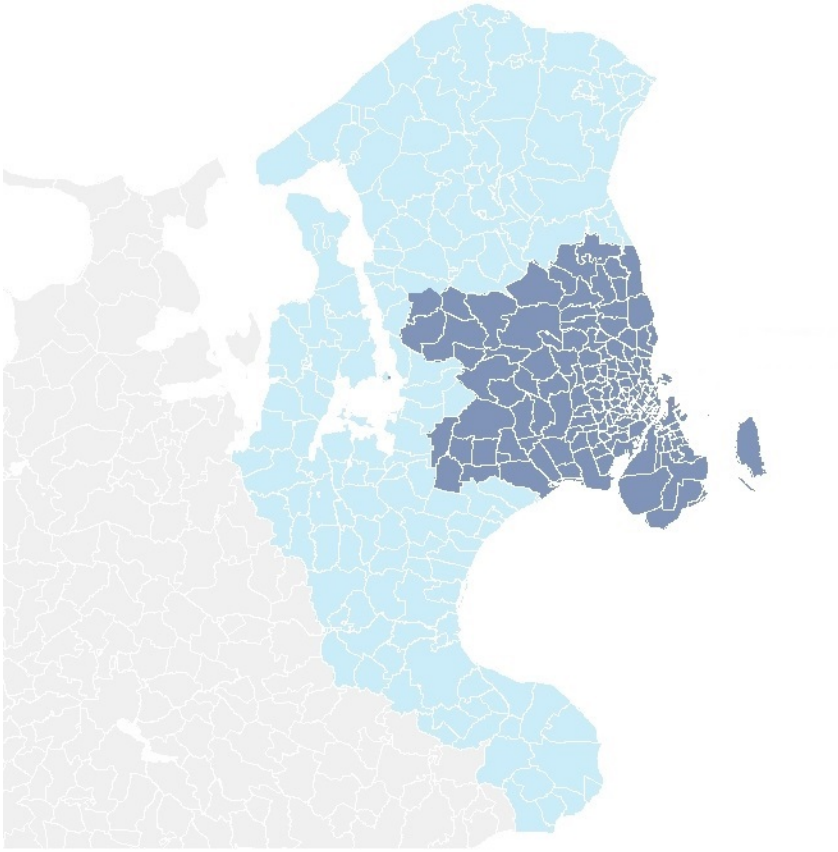
2 Data

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 light blue and dark blue areas of Figure 1 combined and covers 3,038 square kilometers, including even distant suburbs of Copenhagen and substantial areas of open space that are not built up.[5] 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.

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

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.

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. 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.[6] 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.

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

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 above this age who are still living with their parents.[7] 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.

Property Prices For the years 1993 to 2019 we observe the universe of real estate transactions in Denmark including each transaction price and exact location.[8] 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.[9]

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.[10] 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

[7] 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.

[8] Real estate transaction data for Denmark prior to 1993 is unavailable, as digital registration only began in that year. Earlier transactions were recorded manually in physical archives and have not been digitized.

[9] 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.

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

education are low skilled while those with at least a bachelor degree are high skill. 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 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.[11] Panel A of Figure 2 shows the frequency of residential

[11] 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

Figure 2: Residence and Workplace Mobility over the Life Cycle

(a) Probability of moving residence or work



(b) Conditional probability of moving within CMA



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.

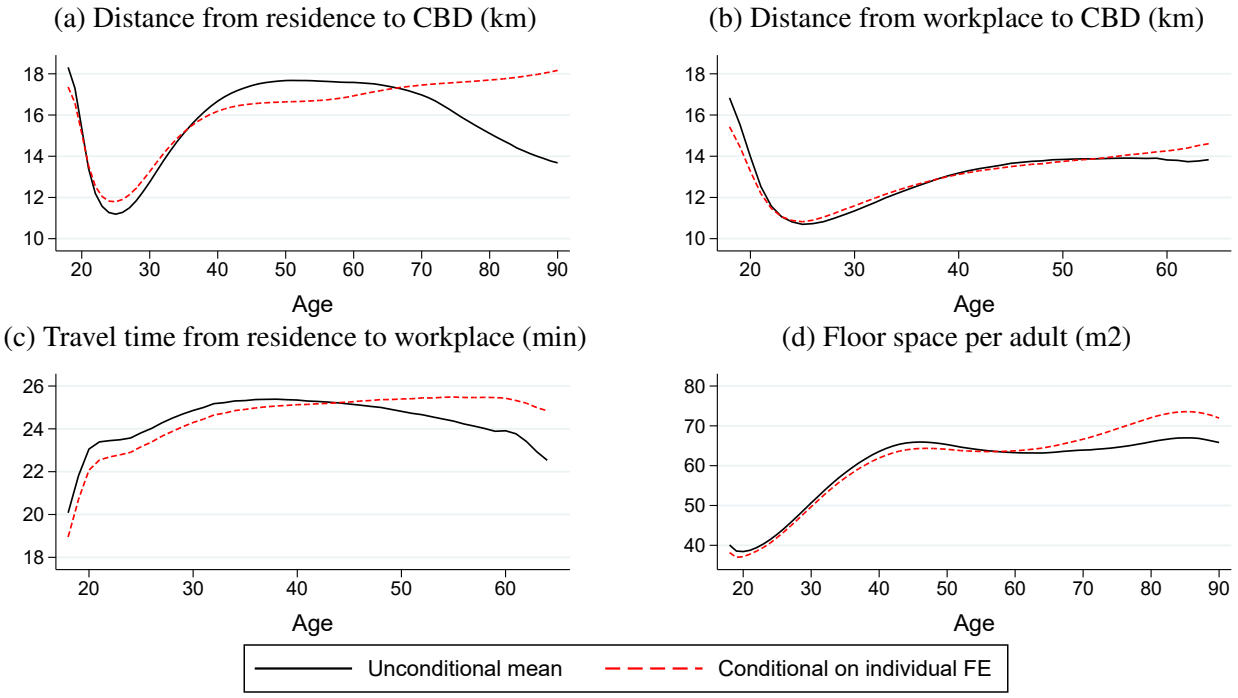
Panel B of Figure 2 explores to what extent people move between different locations within Copenhagen or from Copenhagen to other parts of Denmark or abroad. The figure shows that the vast majority of moves occur within cities, both in terms of residence and in terms of workplace changes. In particular, 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 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 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 they share a household with. We therefore compute a measure of floor space per adult in square meters.

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 year or those who move such short distances that both their origin and destination are in the same 100×100 meter grid cell.

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



In Panel A of Figure 3, we display the average distance in kilometers from of the residence location from 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 are higher ages.

In Panel B of Figure 3, 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 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 Panel C of Figure 3. This patterns 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 Panel D of Figure 3.[12] 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 the 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, when they become widowed, among other important life milestones. For consistency, we define a life event by the order in which they happen (first separation, second cohabitation, etc).[13] 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 B1). 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).[14] Moreover, we consider that some of these events are correlated; individuals

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

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

[14] Online Appendix Figures B6 and B7 show the frequency of the events by age.

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 B1.

We estimate the effects of the multiple life events jointly, using the methodology of Borusyak 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 C. 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).

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.[15] 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 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 layout the necessary assumptions in Online Appendix Section C.

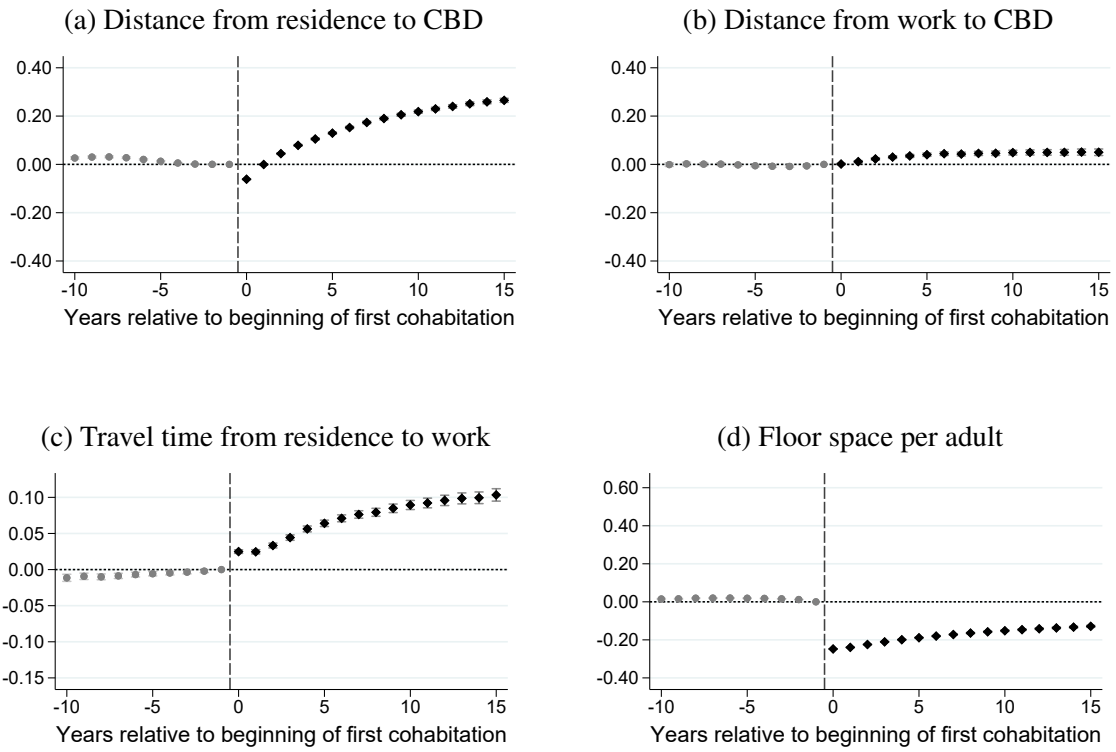
Now we show a simplified and concise version of our method in one equation only. Consider a set of life events $e \in \mathbb{E}$, for which we want to estimate the average treatment effects of. Equation 1 shows the event-study regression in a single step:

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

where Y_{it} is the outcome of individual i in year t. The estimated individual and age fixed effects are represented by $\hat{\theta}_i$ and $\hat{\eta}_a$, respectively. $K_{it}^e = t - E_i^e$ 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

[15] 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



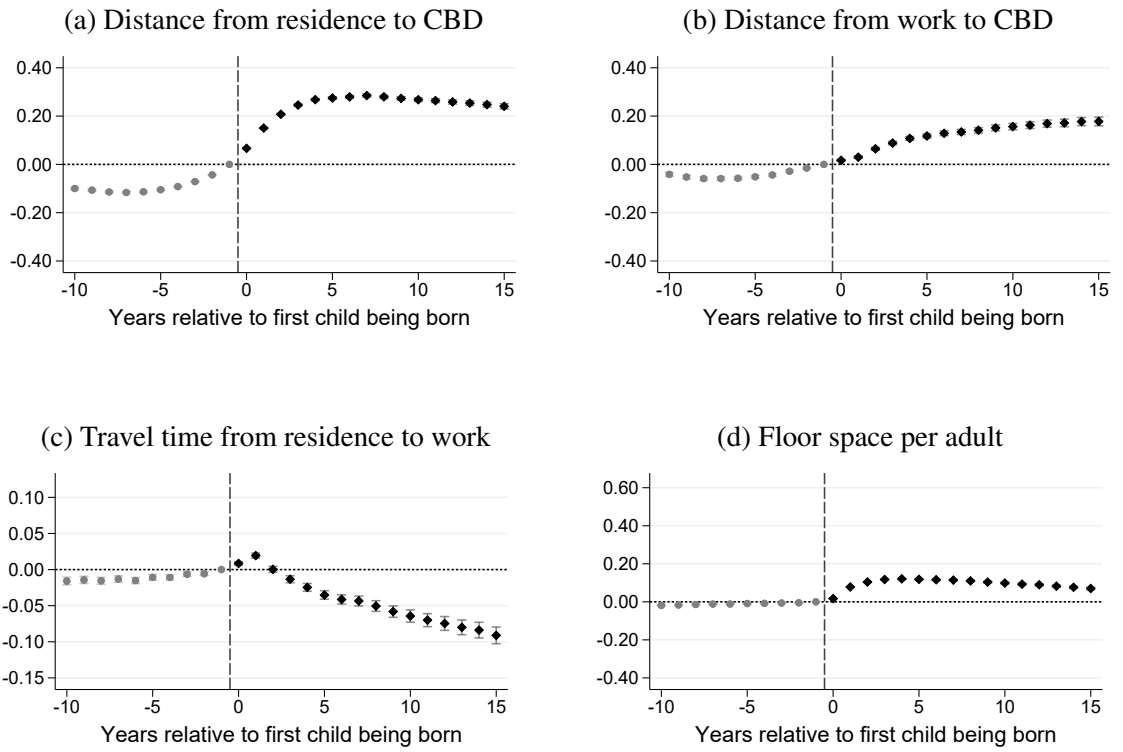
effects from -10 to $+15$. We consider two different sets of events $\mathbb{E}$ for our regressions, early and late life events. Implicitly, we assume that these sets of events are uncorrelated, which, considering the distance in ages between them this seems a reasonable assumption.

We look at the effects of life events on four main outcomes: distance from residence from 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.[16] The figure plots the point estimates and the confidence intervals, but due to the large sample size the standard errors are barely visible. Panel A 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. Panel B 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 Panel A. Consequently, we observe a relative increase in travel time between residence

[16]The sample used to estimate the early life events contains 1,323,393 unique individuals. Approximately 75% of the individuals cohabit before their first child, while 12.6% experience both in the same year. Figure B5 shows how the majority of our sample experiences first cohabitation before having a first child.

Figure 5: Effects of First Child

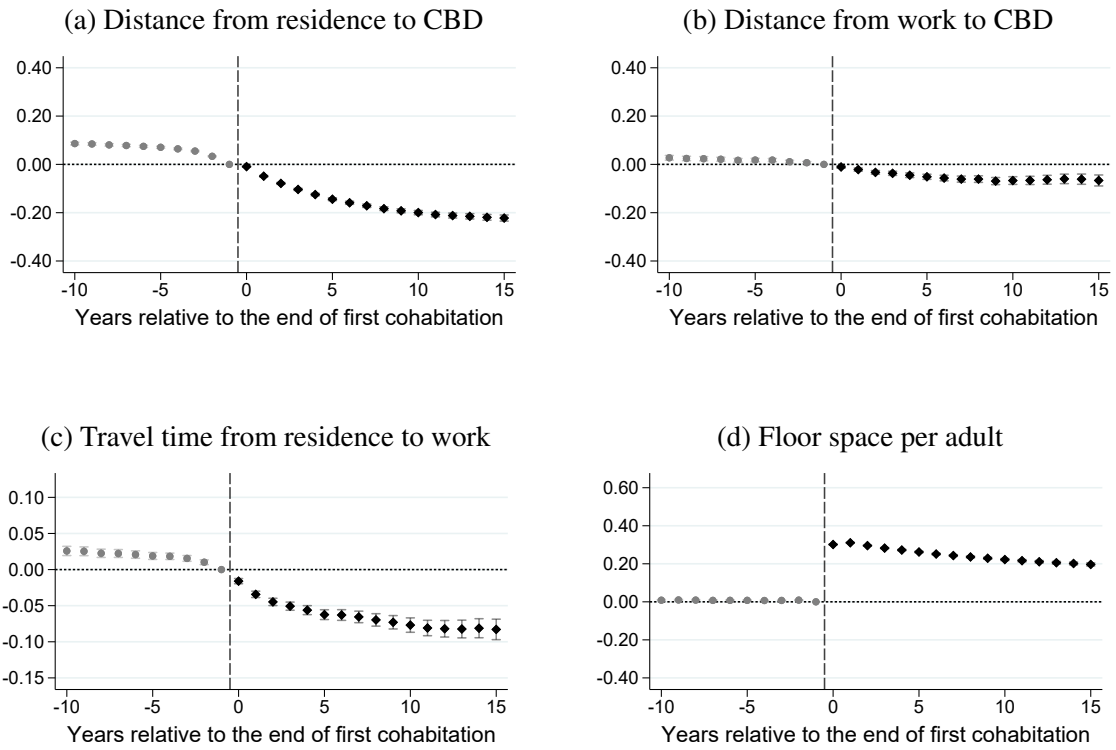


and workplace. Notably, Panel D shows that when individuals start cohabiting, they decrease their consumption of floor space per adult by about 25%, which illustrates that cohabitation rates are an important driver of residential floor space demand in cities.

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.[17] The two top panels suggest that there are some anticipation effects with respect to both residence and workplace location of having a child. Panel A shows that households suburbanize after the first born arrives, and this effect is large and persists for many years after the child is born. Individuals also move their workplace location in response to the first child, with their employment location becoming less central. Panel C 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 the following years. Panel D shows how increasing 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

[17] 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



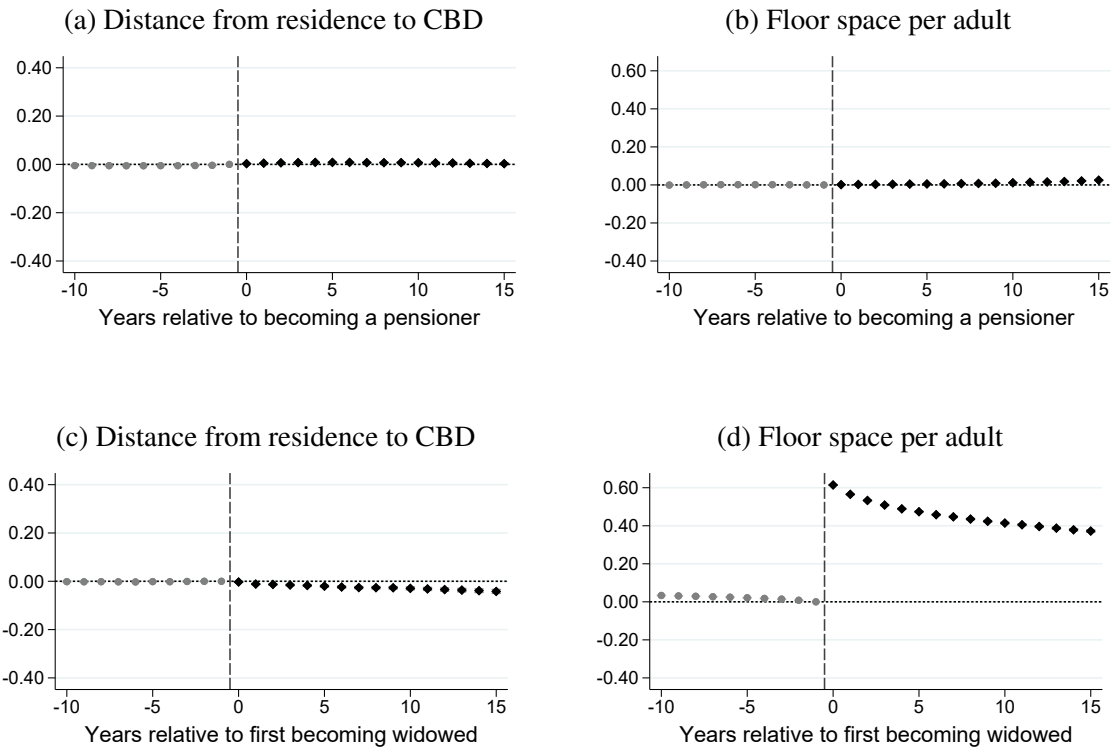
do not see any anticipation effects in floor space consumption.

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. Panel A indicates that individuals slightly anticipate separation, and start moving more centrally even before they decide to stop living together. Panel B shows that the individuals respond by changing workplace right after getting separated, choosing on average more central employment locations. Panel C shows a decline in commuting time after a separation, a result of centralizing both residence and workplace. After separating, 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 by due to different 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.[18] 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 B1, we observe that while 49.6% of the sample experiences a first separation, only 16.9% experience a second.

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

Figure 7: Effects of Retirement and Widowhood



We now present the results with the treatment effects of the late life events, which are estimated separately from the early life events.[19] 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 don't 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 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 Panels A and B, 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, Panel C 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, Panel D 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

[19] See Online Appendix C for the details of our methodology.

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 the all the life events:[20]

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

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:

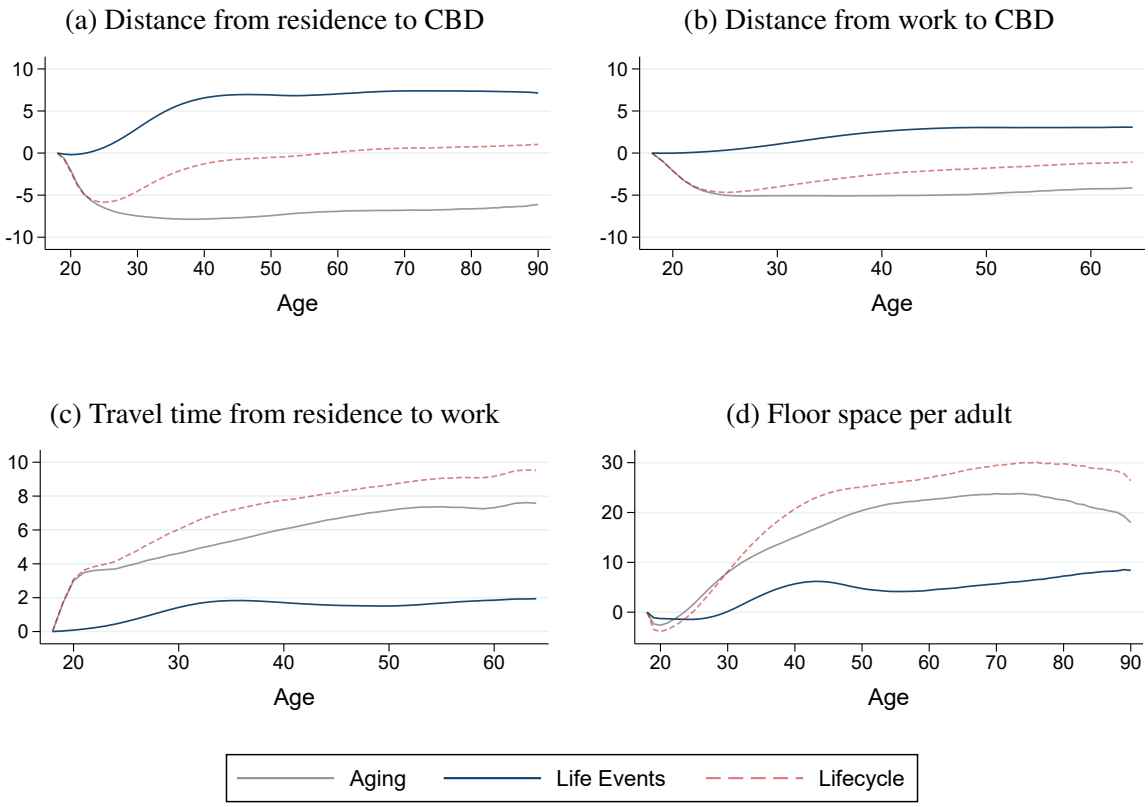
$$\bar{\mathcal{A}}_s = \bar{\mathcal{Y}}_s - \bar{\mathcal{L}}_s \tag{5}$$

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 $\bar{\mathcal{A}}_s$. Panel A decomposes the distance from residence to the CBD, and it shows that the pattern of centralization is counteracted by the force of

[20] 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.

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



life events, pulling individuals to 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 Panel B of Figure 8, we observe two patterns: i) life events have a relatively smaller contribution to where you work to aging, and ii) in the absence of life events, individuals barely decentralize their workplace. The results of Panel C summarize the previous two results. We find that the life cycle in the travel time from residence to home is mostly composed of aging affects, but it is incremented 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. 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 the 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 over the life cycle.

4 Theoretical Framework

We develop a quantitative spatial model in which groups of people who differ in terms of their age and socioeconomic characteristics sort across locations in the city. The model builds closely on the literature following Ahlfeldt et al. (2015) and differs from the existing literature in two main ways. First, we consider both workers and non-working population such as pensioners and students, which account for substantial parts of city populations. Second, due to the richness of our data, we can model a much larger number of different groups of workers and the non-working population 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 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.[21]

To capture the rich heterogeneity of behavior revealed by the stylized facts in Section 3, the model allows agents to differ along three key dimensions: life stage (students, young worker, old worker, and pensioner); skill level (high or low); and family type (single, cohabiting and cohabiting with children). This results in a total of 12 types of workers and 5 types of non-working agents.

4.1 Workers

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 family types of workers f, which are determined by marital status, children 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

[21] 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.

common 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.[22]

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 has 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 in location n is determined by:

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

4.2 Non-Workers

The non-working population in our model consists of pensioners and students, who only make a choice about their place of residence and do not have a workplace.[23] While we do not further disaggregate the group of students, we allow for four different types of pensioners who are defined by skill (high or low skill) and family type (single or cohabiting). This results in a total of five different types of non-working agents.[24]

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

[23] While one could assume that students commute to their university to study, we do not observe the exact location at which students attend lectures and seminars and therefore abstract from this aspect of their life.

[24] Another large group of non-working population are children living with their parents. As such children do not make an independent location decision we assume that the preferences of parents incorporate the utility of their children.

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:

$$U_n^r(\rho) = \frac{B_n^r \bar{w}^r z_n^r(\rho)}{(P_n)^{\alpha^r} (Q_n)^{1-\alpha^r}} \quad 0 < \alpha^r < 1 \tag{9}$$

The idiosyncratic taste shocks $z_n(\rho)$ influence the residential location choice, and are independent drawns from a Fréchet distribution with commulative 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:

$$\lambda_n^r = \frac{L_{Rn}^r}{L^r} = \frac{(B_n^r)^{\varepsilon^r} \left(Q_n^{1-\alpha^r}\right)^{-\varepsilon^r}}{\sum_{k \in \mathbb{N}} \left(B_k^r\right)^{\varepsilon^r} \left(Q_k^{1-\alpha^r}\right)^{-\varepsilon^r}}, \quad n \in \mathbb{N}, \tag{10}$$

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.

4.3 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), but 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 used all four types of labor inputs (L_{Fi}^o) and commercial floor space (H_{Fi}) as inputs:

$$Y_i = A_i \prod_{o \in \mathbb{O}} \left(\frac{L_{Fi}^o}{\beta_i^o}\right)^{\beta_i^o} \left(\frac{H_{Fi}}{\beta^H}\right)^{\beta^H}, \quad 0 < \beta_i^o, \beta^H < 1, \quad \sum_{o \in \mathbb{O}} \beta_i^o + \beta^H = 1, \tag{11}$$

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:

$$A_i \prod_{o \in \mathbb{O}} \left(\frac{1}{w_i^o}\right)^{\beta_i^o} \left(\frac{1}{q_i}\right)^{\beta^H} = 1, \quad 0 < \beta_i^o, \beta^H < 1, \quad \sum_{o \in \mathbb{G}} \beta_i^o + \beta^H = 1, \tag{12}$$

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

4.4 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.

$$L_{Fi}^{of} = \sum_{n \in \mathbb{N}} \lambda_{ni|n}^{of} L_{Rn}^{of} \tag{13}$$

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:

$$\lambda_{ni|n}^{of} = \frac{\left(\mathcal{B}_i^{of} w_i^o / \kappa_{ni}^{of}\right)^{\varepsilon^{of}}}{\sum_{\ell \in \mathbb{N}} \left(\mathcal{B}_\ell^{of} w_\ell^o / \kappa_{n\ell}^{of}\right)^{\varepsilon^{of}}} \tag{14}$$

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 $\lambda_{ni|n}^{of}$:

$$v_n^{of} = \sum_{i \in \mathbb{N}} \lambda_{ni|n}^{of} w_i^o \tag{15}$$

4.5 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:

$$H_{Ri} = \sum_{o \in \mathbb{O}} \sum_{f \in \mathbb{F}} (1 - \alpha^{of}) \frac{v_i^{of} L_{Ri}^{of}}{Q_i} + \sum_{r \in \mathbb{R}} (1 - \alpha^r) \frac{\bar{w}^r L_{Ri}^r}{Q_i} \tag{16}$$

$$H_{Fi} = \beta_H \left(\frac{A_i}{q_i}\right)^{\frac{1}{1-\beta_H}} \prod_{o \in \mathbb{O}} \left(\frac{L_{Fi}^o}{\beta_i^o}\right)^{\frac{\beta_i^o}{1-\beta_H}} \tag{17}$$

4.6 General Equilibrium

Given the 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 (10); profit maximization and zero profit condition (12); commuter market clearing for each worker type (13); residential floor space clearing (16); commercial floor space clearing (17). The conditions for the existence of a uniqueness equilibrium of the model are detailed in Online Appendix Section D.

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 striking 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 is to determine the full general equilibrium contribution of each model perimeter for the observed sorting in the city.

Table 1: Estimated Model Parameters by Group

Model Group	Housing Expenditure Share	ν (PPML)	ν (PPML IV)	ε	ϕ
Young, single, low-skill	31.9%	−0.080	−0.127	5.594	−0.023
Young, single, high-skill	30.1%	−0.056	−0.097	6.043	−0.016
Young, cohabiting, low-skill	23.9%	−0.064	−0.104	5.817	−0.018
Young, cohabiting, high-skill	22.3%	−0.043	−0.079	7.862	−0.010
Young, cohabiting with children, low-skill	23.8%	−0.077	−0.122	9.749	−0.013
Young, cohabiting with children, high-skill	23.1%	−0.060	−0.103	8.042	−0.013
Senior, single, low-skill	31.4%	−0.081	−0.130	5.308	−0.024
Senior, single, high-skill	30.6%	−0.064	−0.106	6.000	−0.018
Senior, cohabiting, low-skill	23.9%	−0.079	−0.127	5.431	−0.023
Senior, cohabiting, high-skill	22.2%	−0.066	−0.109	5.455	−0.020
Senior, cohabiting with children, low-skill	23.7%	−0.076	−0.121	7.571	−0.016
Senior, cohabiting with children, high-skill	22.8%	−0.062	−0.101	5.853	−0.017
Pensioner, single, low-skill	37.0%				
Pensioner, single, high-skill	33.4%				
Pensioner, cohabiting, low-skill	31.9%				
Pensioner, cohabiting, high-skill	26.4%				
Student	39.2%				

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.

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 17 groups in the model. We estimate the housing expenditure shares using our data on individual income, the size of residential units and estimated square meter prices. In particular, we calculate the estimated value of a households residential unit and convert this to an annual fair market rent using an assumed rental return. This estimated rental return is divided by the income of all members of the household to arrive at the housing expenditure share of this household. 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 nearly 40% for students, 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:

$$\log \lambda_{ni}^{of} = \gamma_n^{of} + \gamma_i^{of} - \nu^{of} \times t_{ni}^{of} + u_{ni}^{of} \tag{18}$$

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 at an instrument for travel times, which avoids the potential endogeneity in the placement of infrastructure where exogenous travel demand shifters induce planners to invest into 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 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 the presence of worker and non-worker groups and fewer labor market dimensions (o) than worker groups (of). Our departure point is Equation 13, 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

[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.

the commuter market clearing condition as follows, where the only unknown variable is the transformed wages, which are identified up to scale:

$$L_{Fi}^{of} = \sum_{n \in \mathbb{N}} \frac{\omega_i^{of} / t_{ni}^{\nu^{of}}}{\sum_{\ell \in \mathbb{N}} \omega_\ell^{of} / t_{n\ell}^{\nu^{of}}} L_{Rn}^{of} \tag{19}$$

Using employment and resident population data, we use an algorithm to find the transformed wages that satisfy Equation 19. 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] For each group, we estimate ε^{of} through an algorithm that, given transformed wages, it 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. We don't directly estimate ε^r of the non-worker groups, and instead impose that they will be the same as their closest worker group, e.g. a high-skilled cohabiting pensioner will take the same value as a high-skilled cohabiting senior worker.

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 with 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, senior, 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

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

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 $(1-\beta^H)$:

$$\beta_i^o = \frac{\sum_f w_i^o L_i^{of}}{\sum_o \sum_f w_i^o L_i^{of}} (1-\beta^H) \tag{20}$$

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 (12) to recover location productivities A_i. We estimate estimate the group-specific residence amenities for both workers and non workers using the residence location choice probabilities detailed in (8) and (10). 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 workers and non-workers, and the commercial floor space quantities that clear 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 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 cities 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 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.

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

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 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.

To substantiate this intuition Figure 9c 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.

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 4.6.

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 removes difference in amenities across groups by assuming that all groups attach the same amenity value to a location, which we set equal to the weighted average of the amenities of estimated for the location in the baseline where the weights are the relative sizes of the groups in the baseline.

[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 Appendix Figure XX we show the same scatterplot for both parameters individually, and it is clear that the negative relationship is driven by commuting costs by group, and not their Fréchet parameter.

We are currently in the process of having these results disclosed by Statistics Denmark and should have the full results in the next one or two weeks. The introduction describes the patterns of these results in a form that does not violate our disclosure requirements.

6 The Effect of Demographic Changes

We use our quantitative urban model to explore how three major demographic trends — population aging, declining fertility, and rising singlehood — affect the spatial structure of cities. Each of these forces is prominent in many developed economies, and their implications for urban land use and housing markets are poorly understood. Our approach consists of separate counterfactual simulations where we alter the population shares of key groups and solve for the resulting equilibrium distributions of residential population, workplace locations, and rents.

In the first counterfactual, we increase the share of individuals aged 40 and older by 10%. As Panel A of Figure 10 illustrates, population aging generates decentralizing pressures on the urban system. Residential demand shifts away from central locations and into more peripheral and suburban areas. This is consistent with the lower valuation of centrally located amenities among older households.

The effect on employment distribution mirrors the residential shifts, with jobs becoming more dispersed. Floor space prices reflect this outward movement: central areas experience declining or stagnant prices, while some suburban areas experience moderate increases. These patterns highlight how population aging could weaken centripetal forces in urban housing and labor markets and increase suburbanization trends.

The second counterfactual halves the number of families with children, simulating a substantial decline in fertility rates. As shown in Panel B of Figure 10 the reduction in households with children leads to a pronounced re-concentration of the residential population into central urban areas. Families with children tend to prefer larger housing units and more peripheral locations; reducing their share reallocates demand toward central parts of the city.

The employment distribution responds similarly, becoming more concentrated in the urban core, consistent with shorter commutes and a workforce more oriented toward central amenities. Floor space prices increase in central neighborhoods and fall in peripheral areas, indicating a relative increase in the value of central locations in the face of reduced demand for suburban housing.

While each of these demographic shifts generates distinct spatial pressures, their combined effect is less straightforward. On one hand, declining fertility and increasing singlehood both promote centralization, driven by the preferences of smaller, childless households for central amenities. On the other hand, population aging induces decentralization as older households seek more spacious, peripheral residences.

Taken together, these trends create offsetting forces. In cities where population aging dominates, decentralization effects may prevail. In contrast, cities undergoing sharp fertility declines and rising singlehood may see greater re-concentration of residential and workplace demand into urban cores. However, in settings where these trends occur simultaneously - as is common across much of Europe and other high-income regions - the spatial impacts may partially neutralize, generating more muted aggregate effects on urban structure.

We will present the full quantitative results for more realistic counterfactuals following final approval from Statistics Denmark.

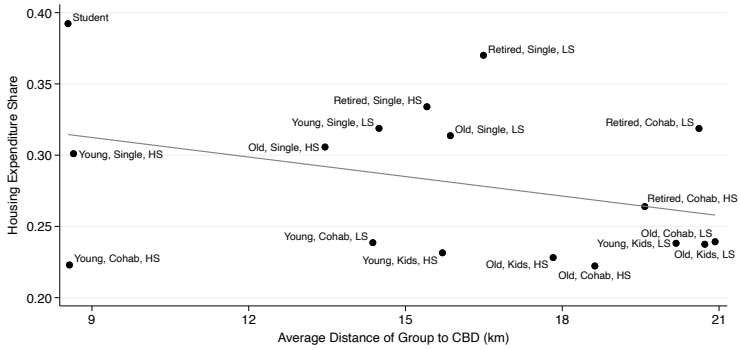
7 Conclusion

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.

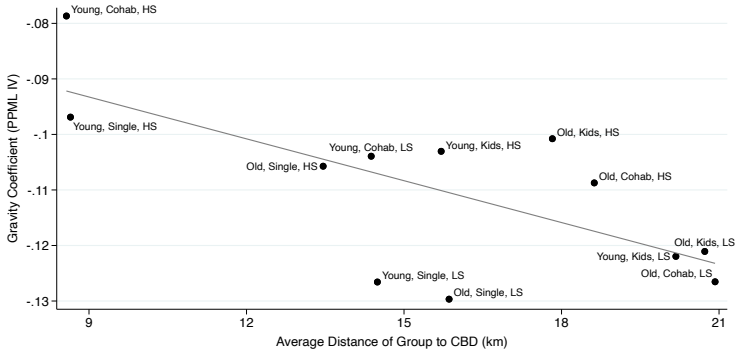
Such shifts in demand caused by demographic change present potential challenges for urban infrastructure and housing markets, as changes in population composition can alter the demand for central versus peripheral locations. However, our counterfactual analysis offers a more nuanced perspective. We show 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. As many cities are exposed to multiple demographic trends simultaneously, these opposing forces may partially offset each other.

This interaction suggests a degree of resilience in urban spatial structures to demographic change. Nevertheless, policymakers should remain attentive to localized pressures, particularly where one trend dominates. Our quantitative approach is easily transferable to other cities, their spatial structures, and anticipated demographic trends, making it a convenient tool for spatial planning at times of rapid demographic change.

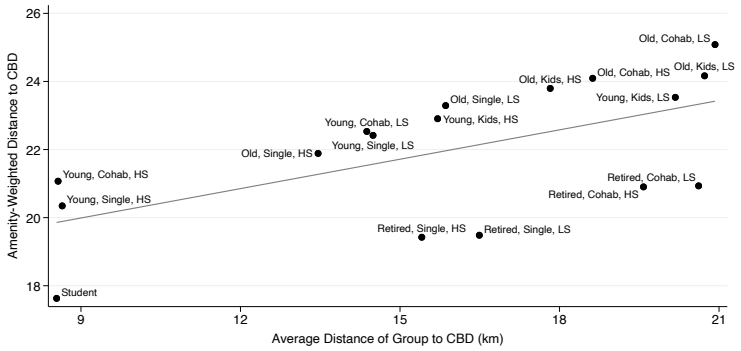
Figure 9: Determinants of spatial sorting



(a) Housing Expenditure



(b) Gravity PPML IV

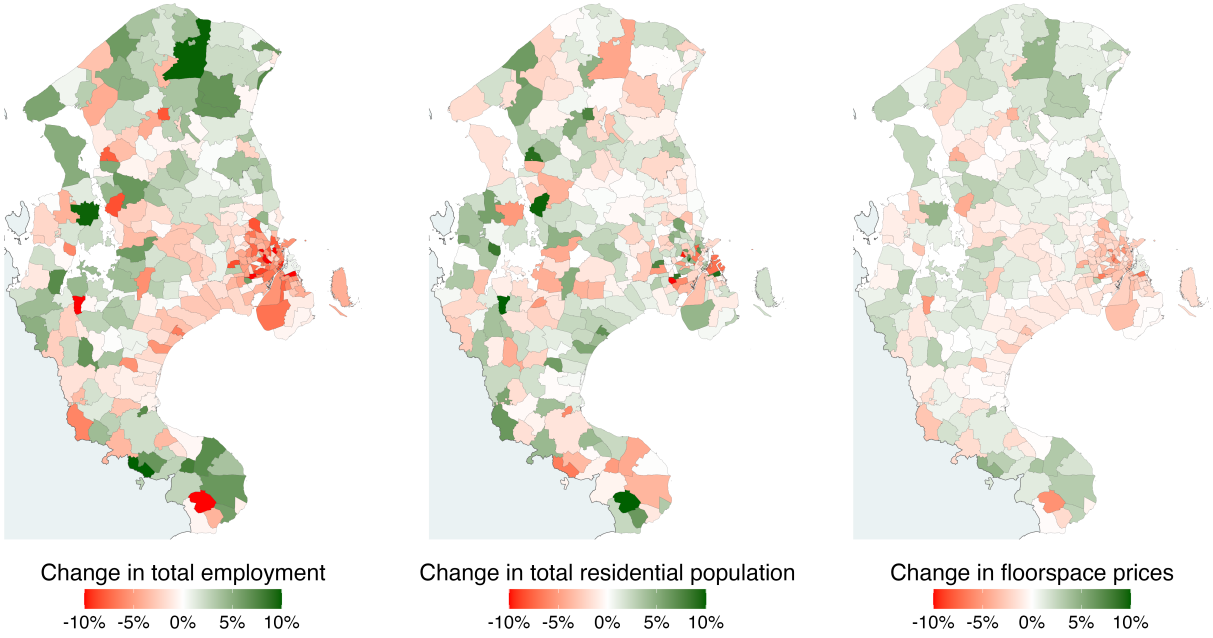


(c) Amenities

Notes: Average distance of group to CBD is the 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.

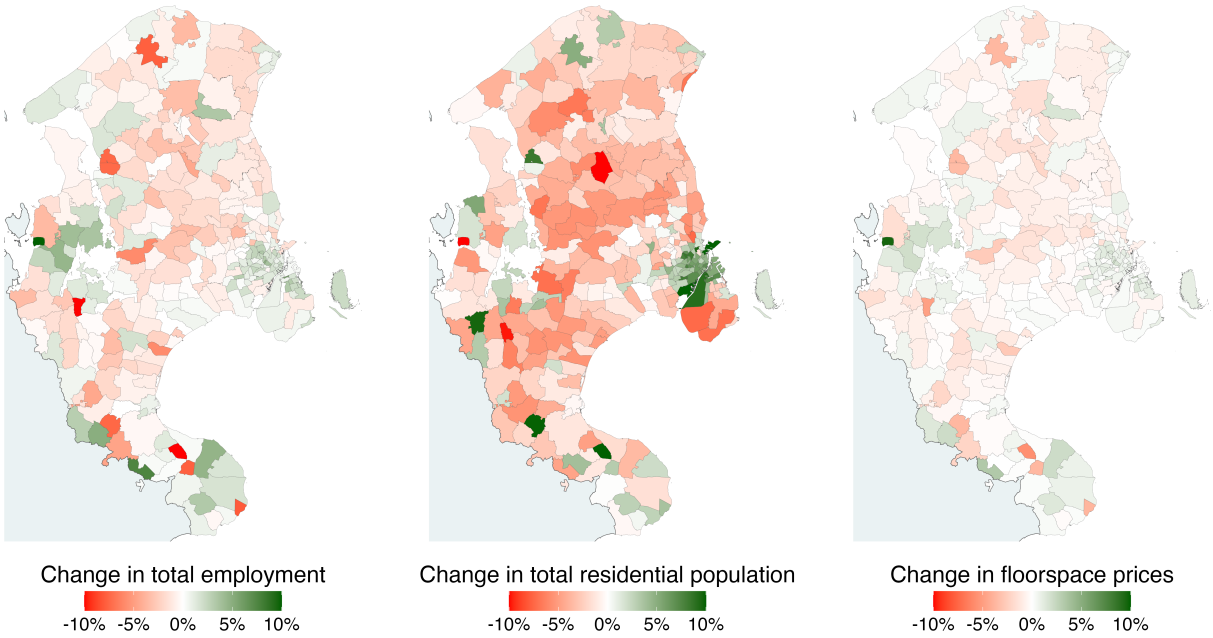
Figure 10: Counterfactual simulations of demographic change

Panel A: Ageing



(a) Emp. pop. total (b) Res. pop. total (c) Res. prices

Panel B: Fertility



(d) Emp. pop. total (e) Res. pop. total (f) Res. prices

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. (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.

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.

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

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.

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.

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.

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.[30] 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[31]. For work characteristics, we use the register IND to obtain information on income, IDAP to record years of work experience,

[30] 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%.

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

IDAN for data on workers wages, occupation and sector. We use the death census (DOD2019) to flag if individuals have died, and the 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 floorspace per adult. Travel time variables are all in minutes, straight line distance variables are in kilometers, and floorspace 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. Floorspace per adult is defined as the total floorspace of the household/adddress/dwelling divided by the number of floorspace-consuming adults. We define a floorspace consuming adult as anybody who is 18 or older that does not live with their parents. We also winsorize the floorspace 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.[32] 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.[33]

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.

[32] 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.

[33] 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.

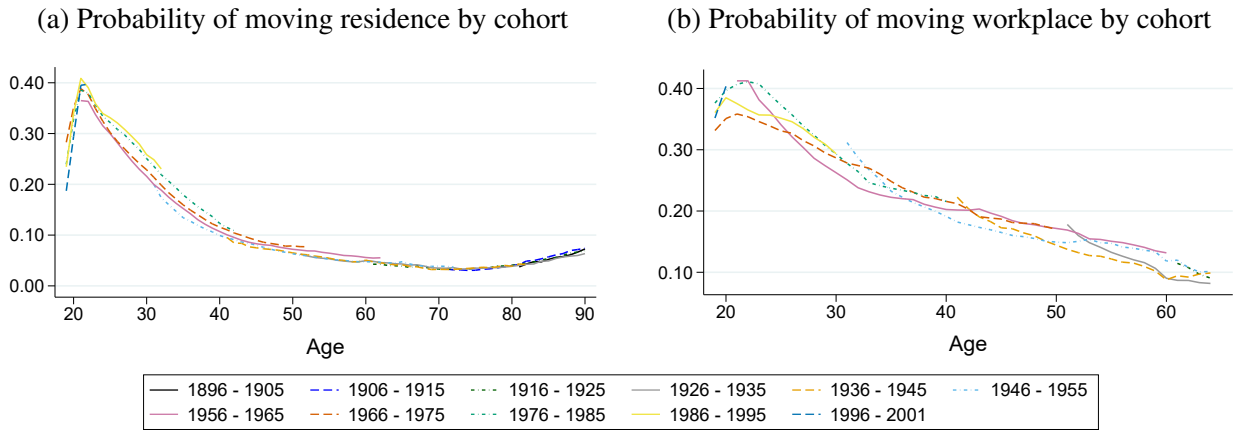
B Stylized Facts

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. Furthermore, we provide a simple econometric framework for our event study methodology and relate it to the existing literature.

B.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 B1 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. Panel A shows a stable life cycle pattern in residential mobility, with close to no differences across cohorts. Panel B 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 B1: Residence and workplace mobility over the life cycle



In Figure B2, we replicate our methodology for all the commuting zones in Denmark.[34] The

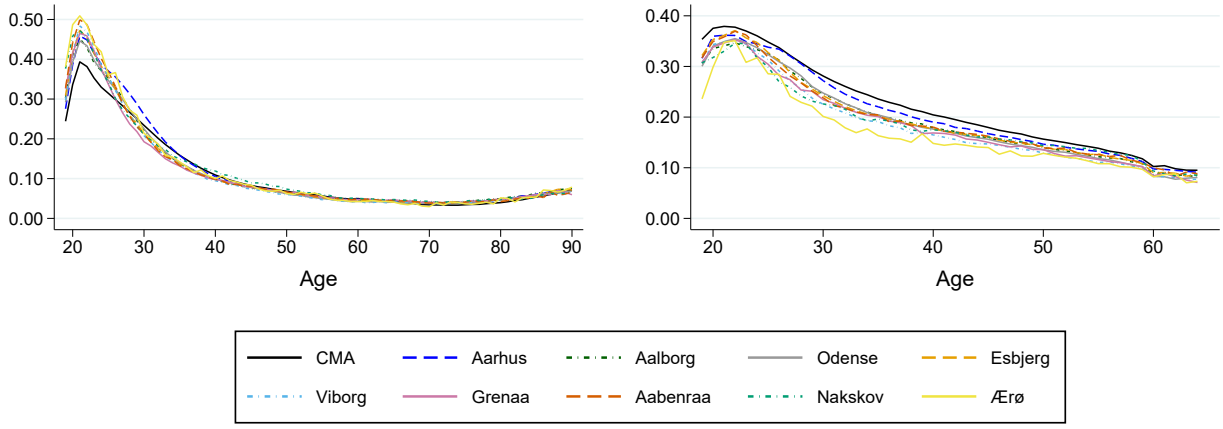
[34]There are XX commuting zones in Denmark, and we consider all but one XXX.

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 B2: 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



B.2 Location Choices over the Life Cycle

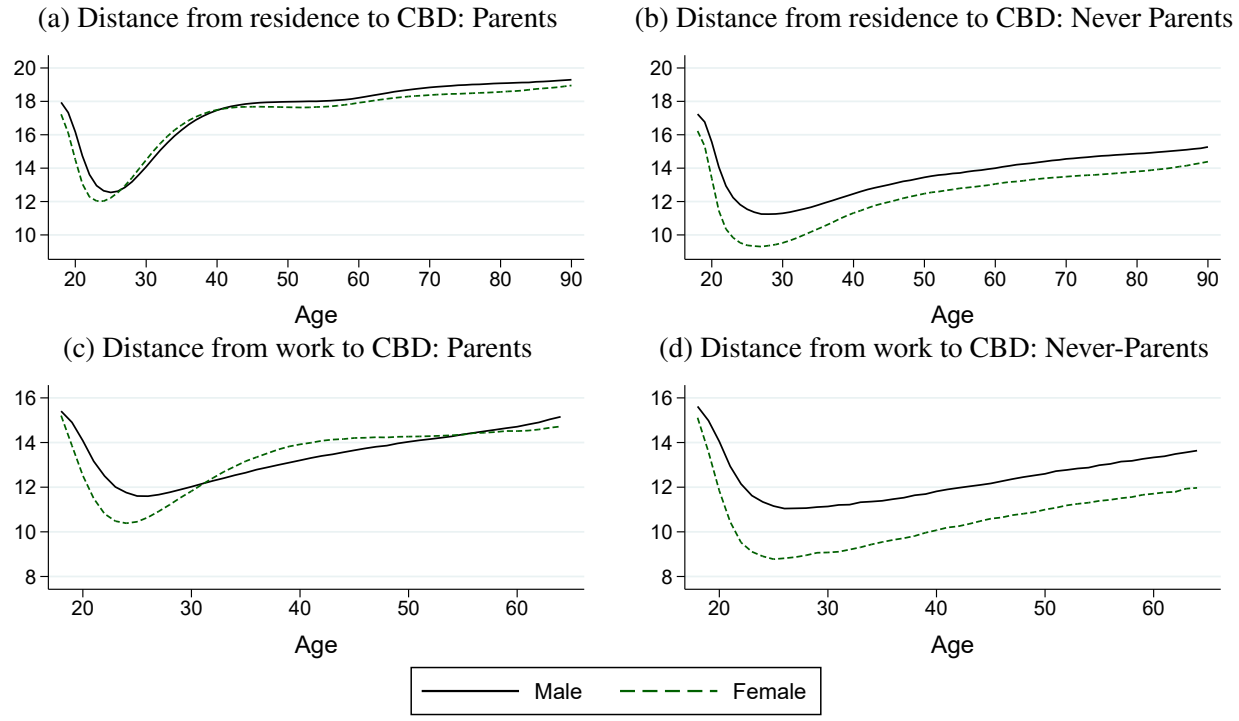
We follow by considering heterogeneities related to the location choices over the life cycle. In Figure B3, we explore if gender and parenthood could be important factors in explaining the patterns observed in Figure 3. In this exercise, we show the mean distances between residence and workplace to the 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 the top two panels of Figure B3, 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.

The two bottom panels of Figure B3 show the mean distance between workplace and CBD,

conditional on individual fixed effects. Comparing Panels C and D, 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.

Figure B3: Residence and workplace location choices over age, by gender and parenthood



In Figure B4, we compute the same outcomes as in Figure 2 conditional on individual fixed effects, by birth decade cohort. The results show striking patterns. Panel A shows that there 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 Panel A of Figure B1. The degree of centralization that occurs in young adulthood seems to vary considerably by cohort as well.

In Panel B of Figure B4, 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. Panel D 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. In Panel C, we verify that commuting time ...

Figure B4: Residence and workplace location choices over age, by cohort



B.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 B1, we list the events we consider with information on the number of treated individuals, their size my

B.3.1 Comparison between our strategy and OLS

B.3.2 Heterogeneities

Gender

Skill

Table B1: Age Distribution of Life Events and Sample Characteristics

Event	Age at Treatment p10	p50	p90	Number of Treated Individuals	Share of Treated Individuals(%)
Early Life Events					
First Child	23	28	35	452,652	58.18
Second Child	26	31	38	298,751	38.40
Third Child	28	34	40	74,029	9.51
First Cohabitation	20	25	33	700,479	90.03
Second Cohabitation	24	30	41	271,156	34.85
Third Cohabitation	28	34	47	83,531	10.74
First Separation	22	27	40	386,092	49.62
Second Separation	26	33	46	131,732	16.93
Late Life Events					
Empty Nest	45	52	60	481,211	20.10
Pension	47	62	67	499,055	20.84
First Widowhood	56	65	79	103,446	4.32

Figure B5: Distance between Year of First Child and Year of First Cohabitation

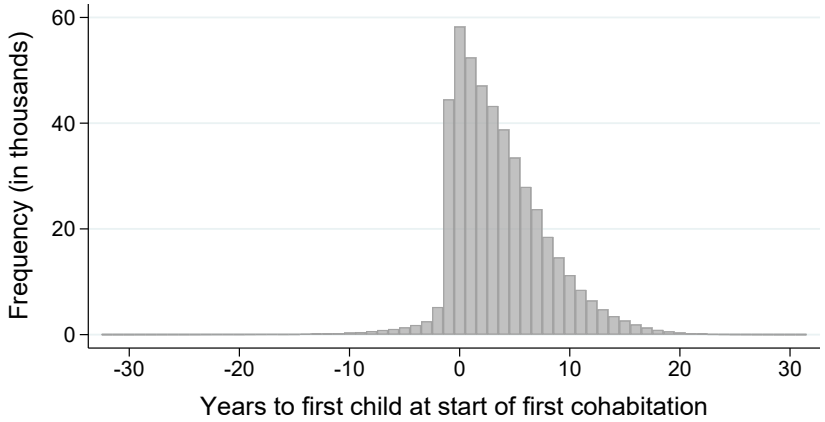


Figure B6: Frequency of Early Life Events by Age

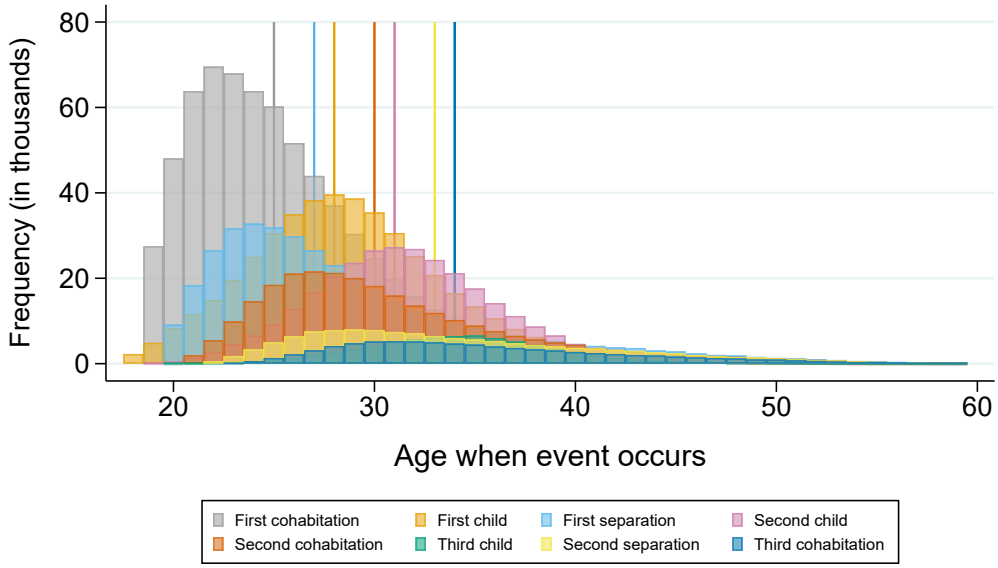
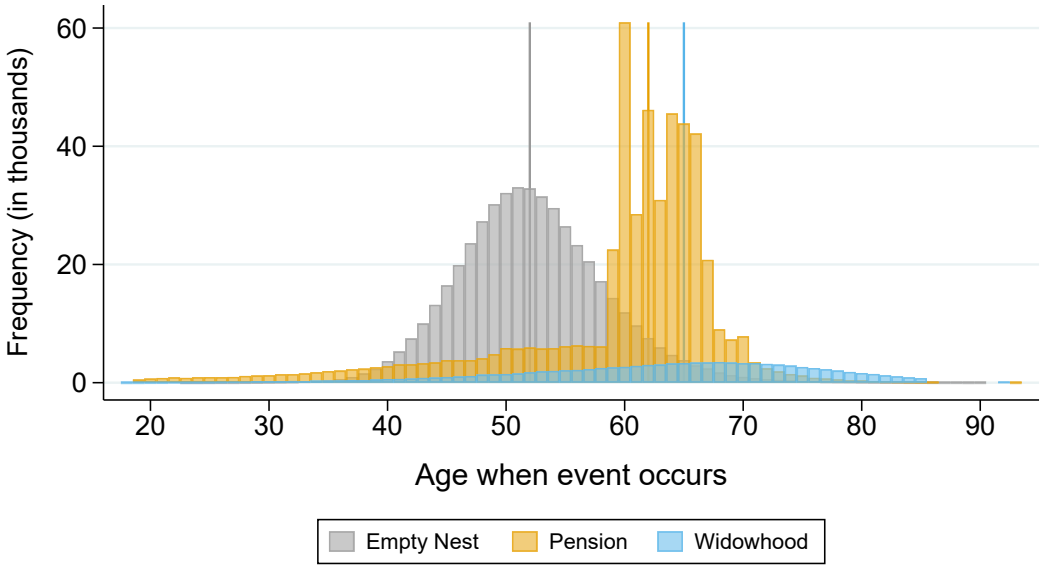


Figure B7: Frequency of Late Life Events by Age



C An Econometric Model of Life Cycle in the City

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 ageing alone?

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

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

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:

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

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:

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

D Quantitative Model

To be completed.

E Model Estimation

E.1 Calibrating the Group Sizes

Before describing the details of the model estimation, it is important to detail how the groups we consider in the model appear in the data. Importantly, we consider three important dimensions to classify individuals: life stage, skill level, and family type. The groups are described in Appendix Table E1. There are two life stages in which individuals do not work: when they are students and when they retire. When they are workers they can be either young workers (age 18 to 45) or senior workers (45 to 65). The state pension age in Denmark is 65, so we choose that as our cut-off to classify pensioners. We often refer to life stages as age groups, since they are closely related to age. Appendix Figure E1 shows the sizes of the groups as a share of the CMA adult population and how stable these are over time.

The second relevant dimension to classify individuals is skill. Students are themselves a category of skill, since they are still in the process of accumulating human capital. The other individuals are divided into low- and high-skill, where high-skilled is someone with at least some tertiary education.[35] Lastly, we classify individuals by the type of family they have: single, cohabiting with a partner, and cohabiting with a partner and children. We do not include single parents in the model since they are a small share of the population as shown in Appendix Figure E2.

Table E1: Relevant dimensions for defining the model groups

	Life stage	Skill	Family type
Non-workers	Students	-	Single
	Pensioners	LS, HS	Single, Cohabiting
Workers	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. We do not consider any heterogeneity for students, so we consider them to have one skill level, and one family type only. Pensioners and workers may have different skill levels and family types. We do not consider pensioners cohabiting with children in the model as they are a very small share of the sample.

Our model considers 17 different groups, out of which 12 participate in the labor market, and 5 don't. These groups account for XX% of the CMA adult population during the period considered, and for XY% of the CMA population in 2015, the year we use to estimate the model.

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

Figure E1: Evolution of age group shares over time

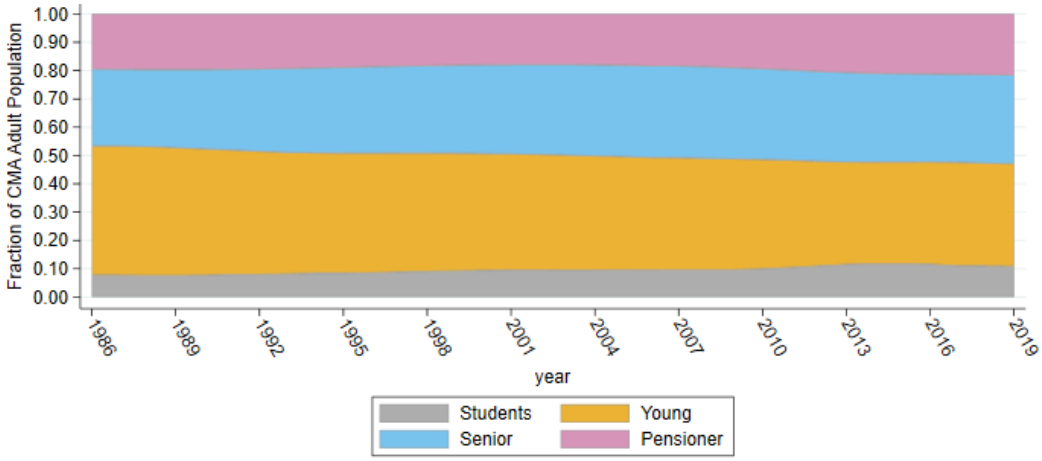


Figure E2: Evolution of family group shares over time

