Online Appendix for "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

July 2025

Table of Contents

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.

A1 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.[1] 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[2]. For work characteristics, we use the register IND to obtain information on income, IDAP to record years of work experience, IDAN

[1] 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%.

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

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

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

A2 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.

A3 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.

A4 Variable Definitions

A4.1 Outcomes

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

A4.2 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.

A4.3 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.[3] 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.[4]

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.

[3] 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.

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

B An Econometric Model of Life Cycle in the City

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

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

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

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

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

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

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

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

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

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

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:

$$\begin{aligned}
\delta^{z1}_{i,a} &= Y_{i,a}(1,0) - Y_{i,a}(0) = \tau^1 \\
\delta^{z2}_{i,a} &= Y_{i,a}(0,1) - Y_{i,a}(0) = \tau^2 \\
\delta^{z3}_{i,a} &= Y_{i,a}(1,1) - Y_{i,a}(0) = \tau^1 + \tau^2
\end{aligned}$$

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 $\hat{\delta}^z_{i,a}$.

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

$$\hat{\delta}^z_{i,a} = \sum_{\ell=1}^{L} d^\ell_{i,a}\tau^\ell + u_{i,a}$$

Assumption 3 is necessary for interpreting the effects of life events as causal. Life events are not randomly allocated, and we don't want to claim that our estimates represent the causal effects of life events on location choices and floor space consumption, as there are other things that may be changing when people choose to go through these life changing decisions, such as a first cohabitation or a first child.[5] Nonetheless, we show (descriptive) evidence that when people choose to experience life events, their location within the city changes, and these results are what matters for our analysis.

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

C Stylized Facts

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

C1 Mobility over the Life Cycle

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

Figure C1: Residence and workplace mobility over the life cycle

(a) Probability of moving residence by cohort (b) Probability of moving workplace by cohort

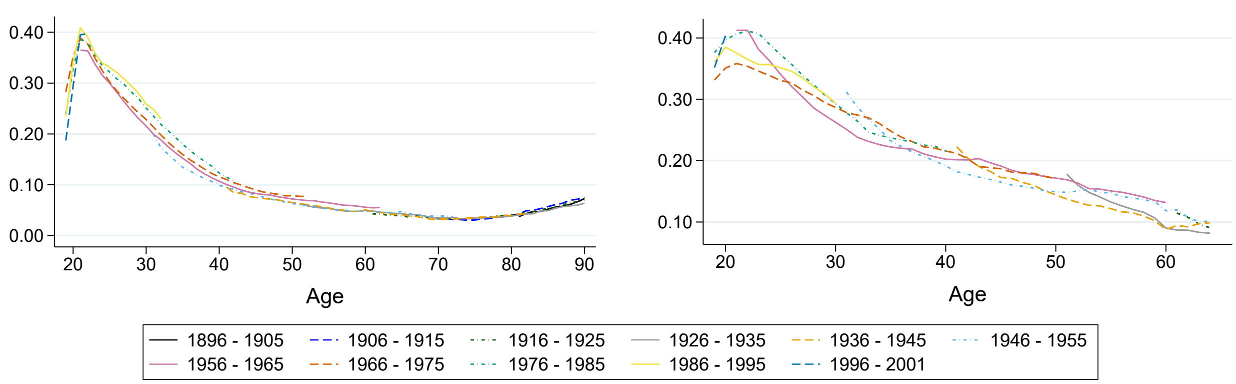

Notes: This figure shows residence and workplace mobility patterns over the life cycle by birth cohort. Panel (a) plots the probability of moving residence against age. Panel (b) graphs the probability of moving workplace against age.

In Figure C2, we replicate our methodology for a subset of commuting zones in Denmark.[6] The

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

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

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

(a) Probability of moving residence by commuting zone

(b) Probability of moving workplace by commuting zone

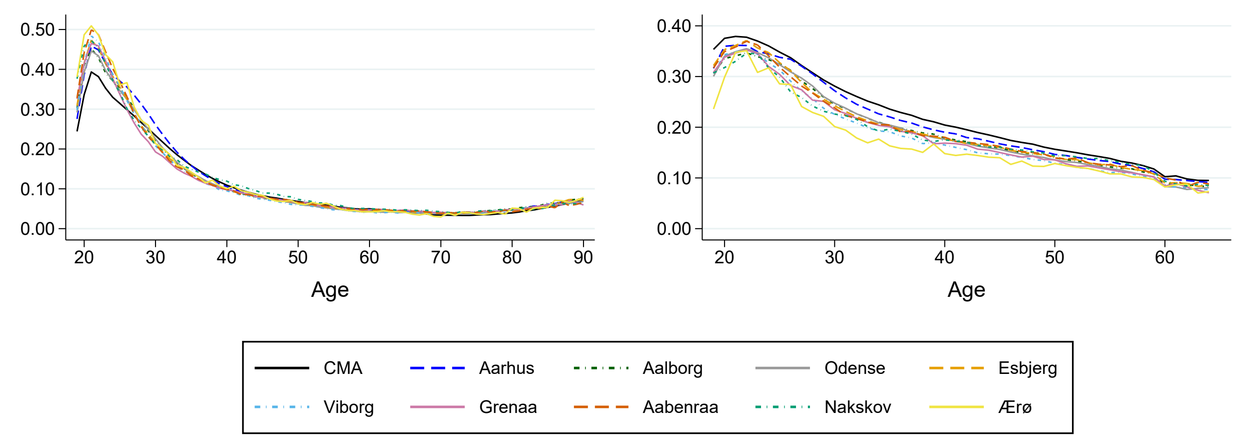

Notes: This figure shows residence and workplace mobility patterns over the life cycle by commuting zone. Panel (a) plots the probability of moving residence against age. Panel (b) graphs the probability of moving workplace against age.

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

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

C2 Location Choices over the Life Cycle

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

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

(a) Probability of moving residence or work across parishes

(b) Conditional probability of moving within CMA

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

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

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

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

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

Figure C4: Distance of Residence and Workplace Moves

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

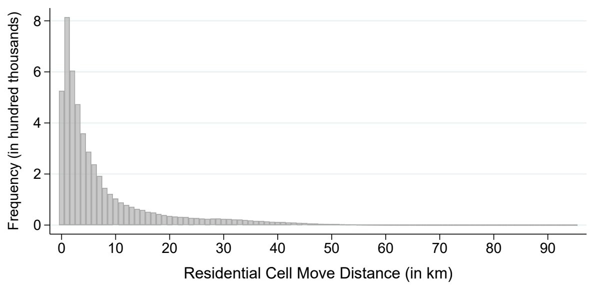

(b) Distance of workplace 100m cell moves (km)

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

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

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

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

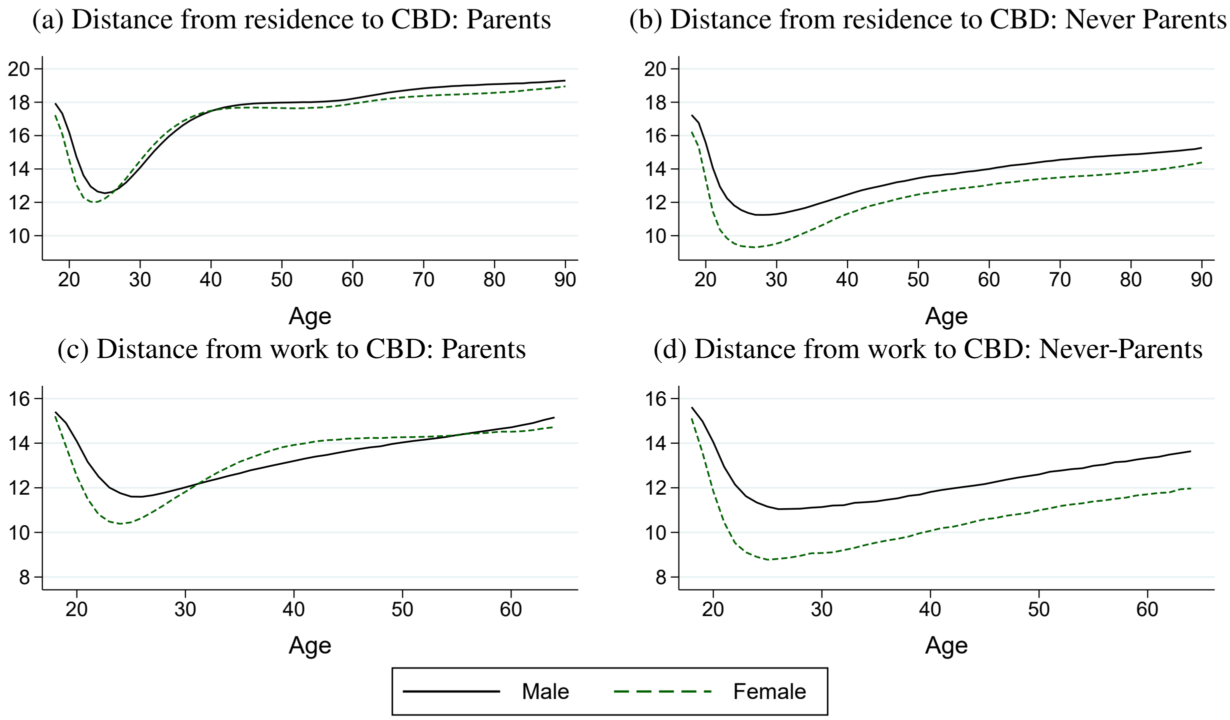

Notes: This figure shows residence and workplace location decisions over the life cycle by gender and parenthood status. Panels (a) and (b) plot the distance in kilometers from residence to the CBD for parents and never-parents, respectively. Panels (c) and (d) graph the distance in kilometers from work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

C3 Life Events

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

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

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

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

(a) Distance from residence to CBD (km) (b) Distance from workplace to CBD (km)

(c) Travel time from residence to workplace (min) (d) Floor space per adult (m2)

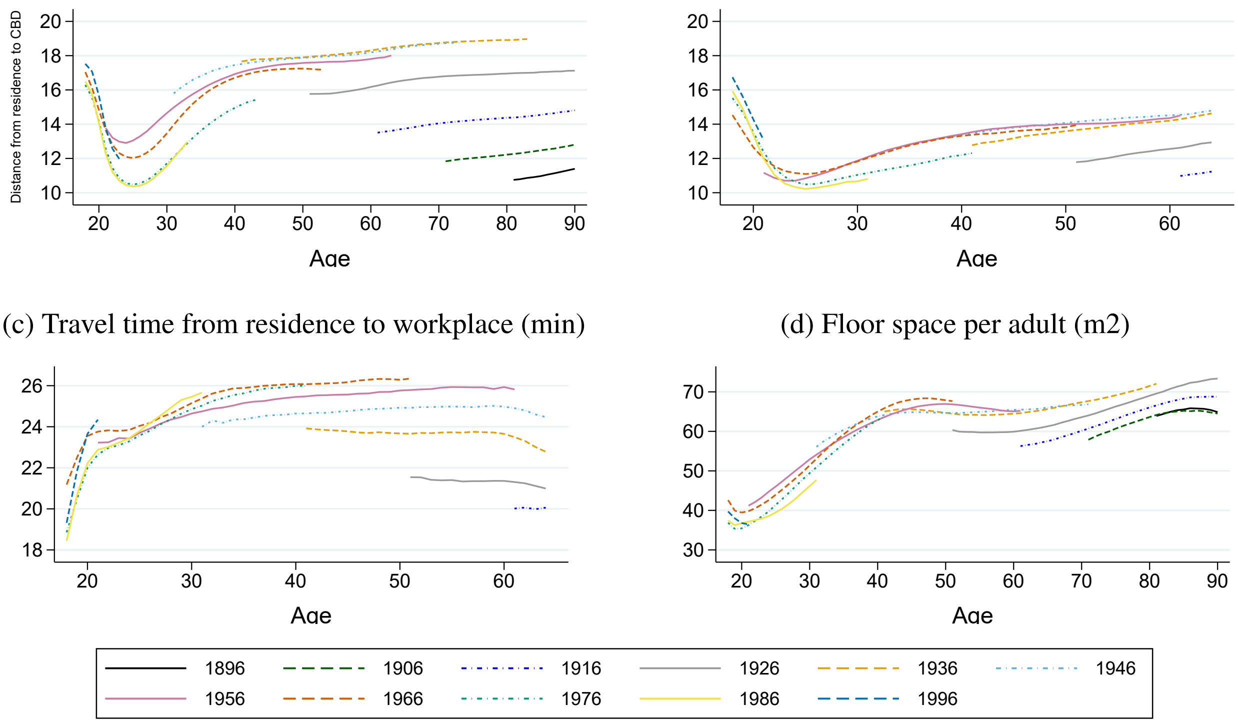

Notes: This figure shows residence and location choices over the life cycle by cohort. Panel (a) plots the kilometer-measured distance from residence to the CBD against age. Panel (b) plots the kilometer-measured distance from workplace to the CBD against age. Panel (c) graphs commuting time in minutes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

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

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

C3.1 Comparison between our Strategy and OLS

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

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

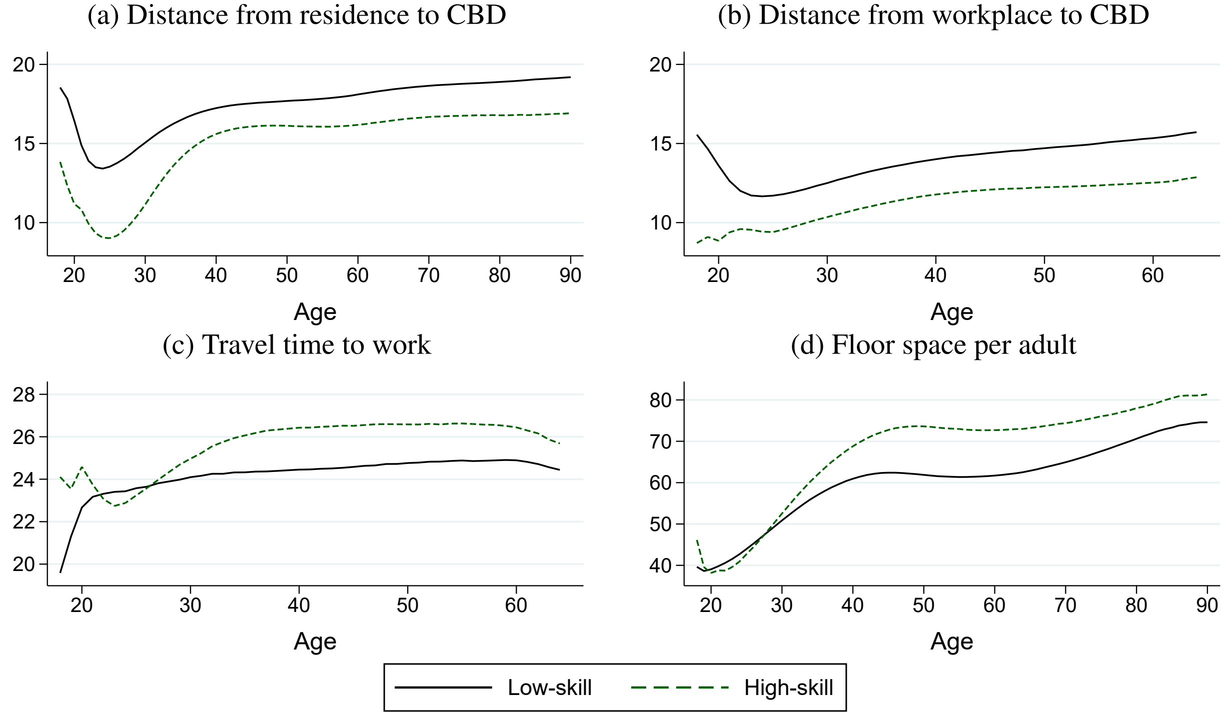

Notes: This figure shows residence and location choices over the life cycle by skill level. Panel (a) plots the kilometer-measured distance from residence to the CBD against age. Panel (b) plots the kilometer-measured distance from workplace to the CBD against age. Panel (c) graphs commuting time in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

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

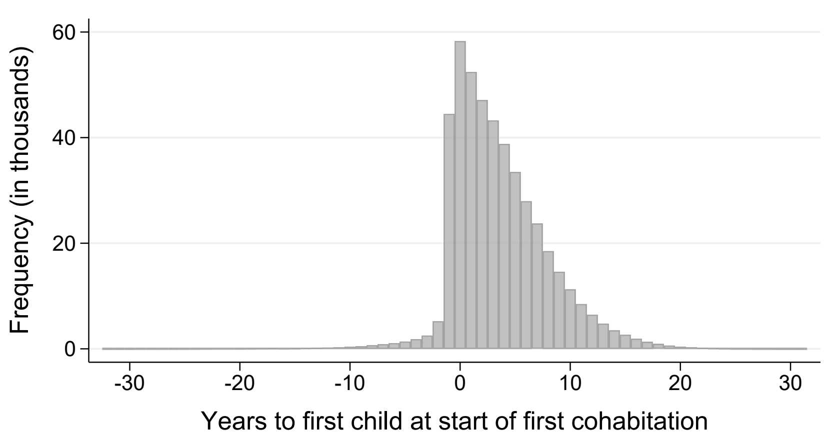

Notes: This figure shows the distribution of time elapsing between the first cohabitation and the birth of the first child.

Table C1: Age Distribution of Life Events and Sample Characteristics

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

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

Figure C9: Frequency of Early Life Events by Age

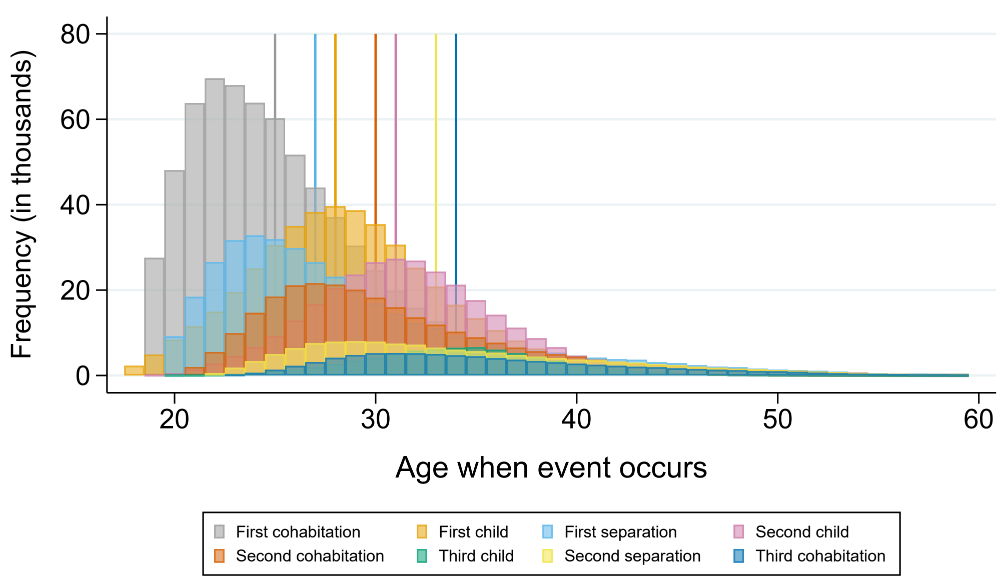

Notes: This figure plots the distribution of early life events against age.

C3.2 Heterogeneities

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

Figure C10: Frequency of Late Life Events by Age

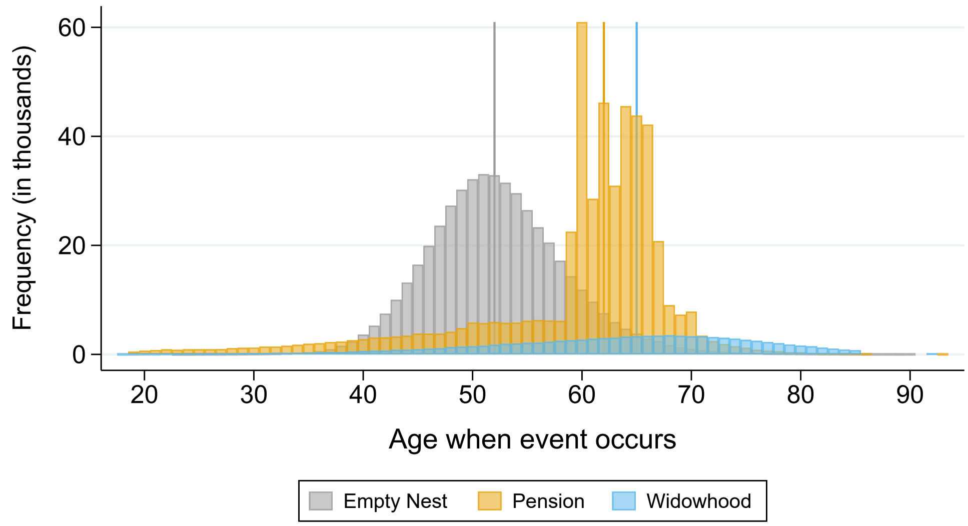

Notes: This figure plots the distribution of late life events against age.

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

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

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

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

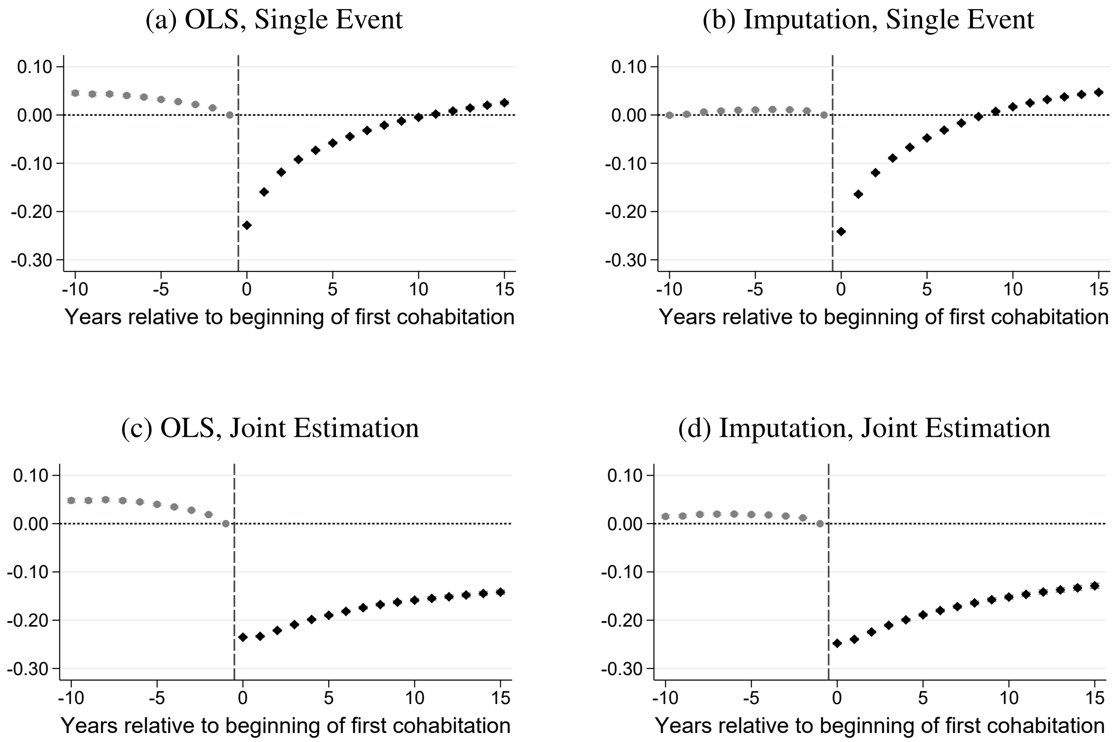

Notes: Each panel of this figure plots event-time coefficients relative to the beginning of the first cohabitation at year 0, comparing OLS to imputation estimation. Panels (a) and (b) show the change in floor space consumption relative to the years from the beginning of the first cohabitation considered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

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

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

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

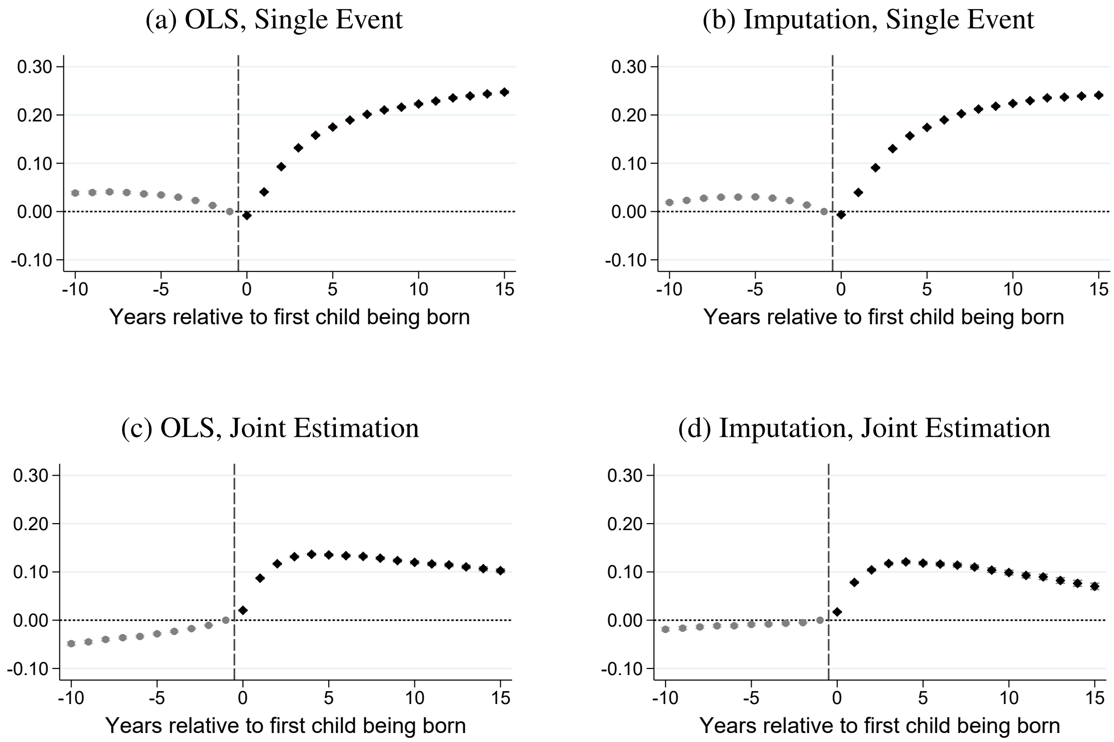

Notes: Each panel of this figure plots event-time coefficients relative to the birth of the first child at year 0, comparing OLS to imputation estimation. Panels (a) and (b) show the change in floor space consumption relative to the years from the beginning of the first cohabitation considered as a single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

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

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

Figure C13: Effects of First Cohabitation: Gender

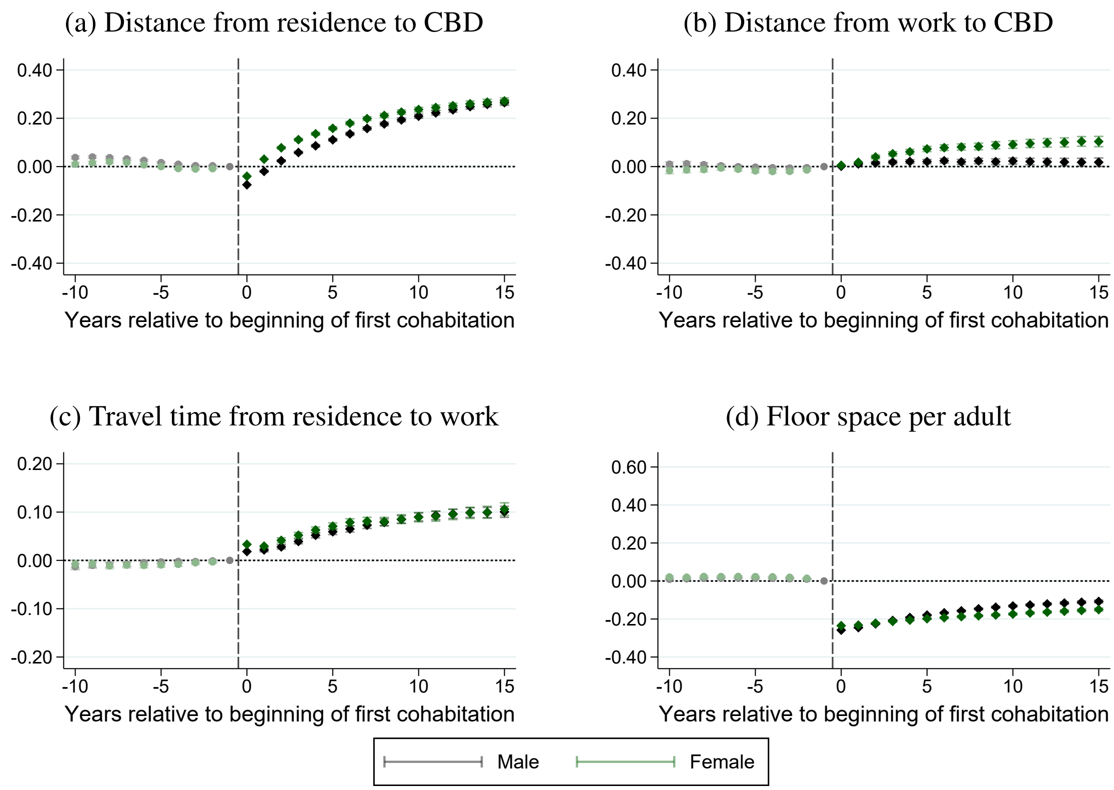

Notes: Each panel of this figure plots event-time coefficients relative to the beginning of the first cohabitation at year 0, disaggregated by gender. Panel (a) shows the change in distance from residence to the CBD, and Panel (b) shows the change in distance from workplace to CBD, both in kilometers.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

Figure C14: Effects of First Child: Gender

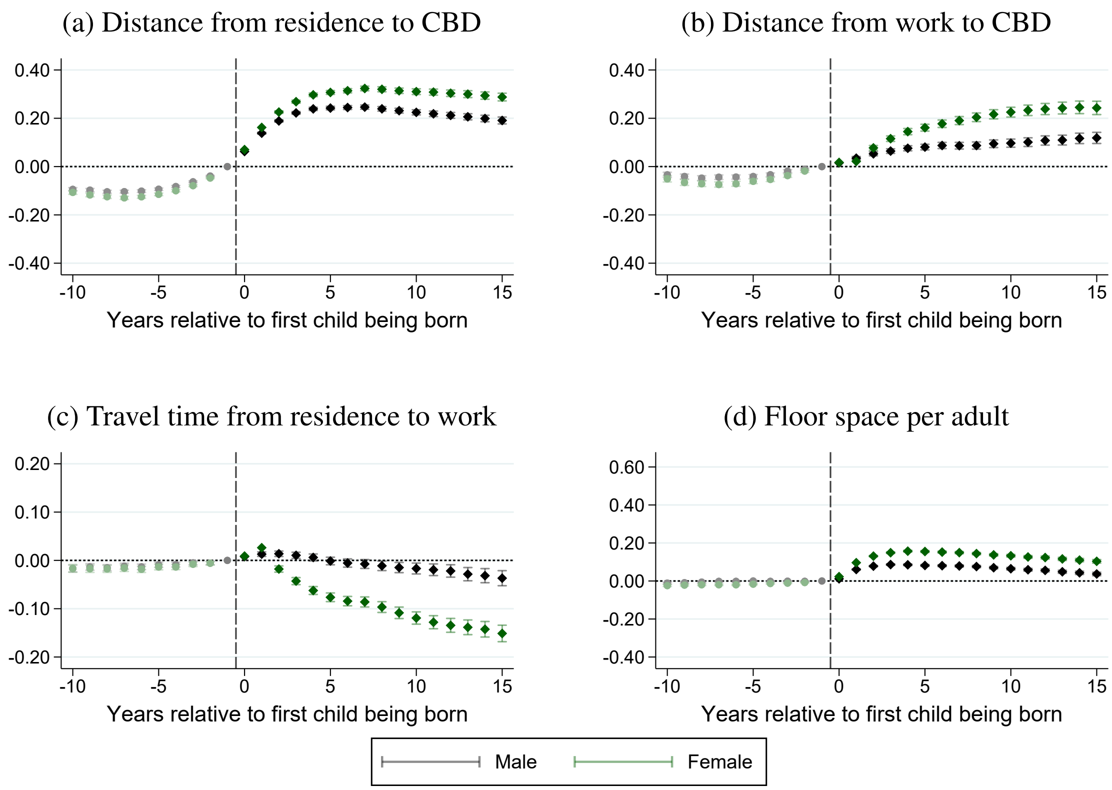

Notes: Each panel of this figure plots event-time coefficients relative to the birth of the first child at year 0, disaggregated by gender. Panel (a) shows the change in distance from residence to the CBD, and Panel (b) shows the change in distance from workplace to CBD, both in kilometers. Panel (c)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

Figure C15: Effects of First Cohabitation: Skill

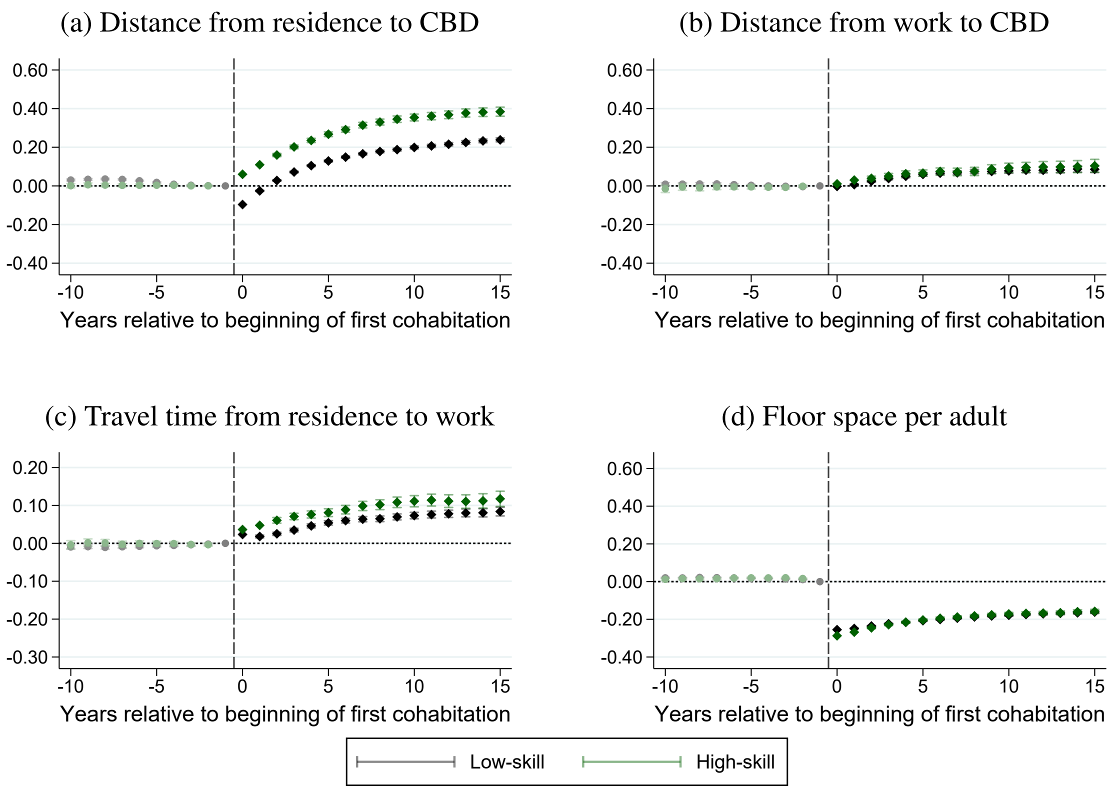

Notes: Each panel of this figure plots event-time coefficients relative to the beginning of the first cohabitation at year 0, disaggregated by skill. Panel (a) shows the change in distance from residence to the CBD, and Panel (b) shows the change in distance from workplace to CBD, both in kilometers.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

Figure C16: Effects of First Child: Skill

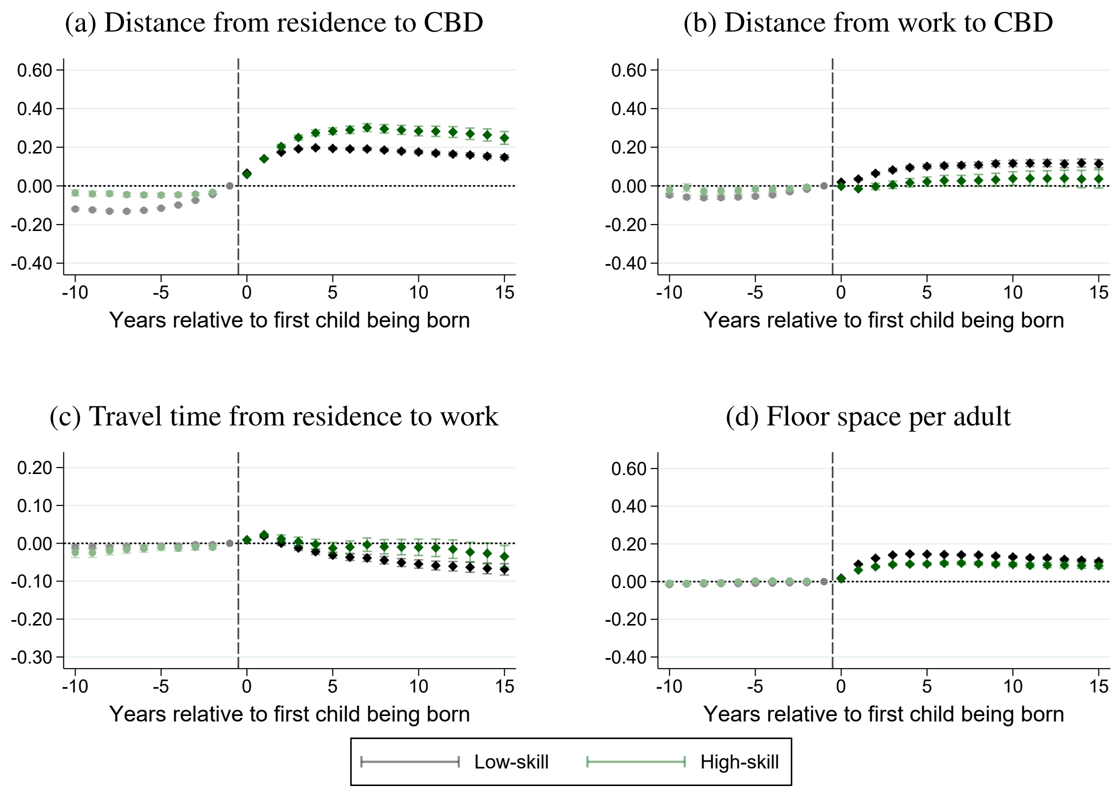

Notes: Each panel of this figure plots event-time coefficients relative to the birth of the first child at year 0, disaggregated by skill. Panel (a) shows the change in distance from residence to the CBD, and Panel (b) shows the change in distance from workplace to CBD, both in kilometers. Panel (c)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

Table D1: Overview of Model Groups in the Extended Model

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

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

D Model

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

D1 Model Extension: Including the Non-Workers

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

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

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

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

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

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

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

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

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

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

E Model Estimation

E1 Calibrating the Group Sizes

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

The second relevant dimension to classify individuals is skill. Individuals are divided into low- and high-skill, where high-skilled is used to describe someone with at least some tertiary education.[7] Students are not considered in this skill classification and excluded from the model population. As shown in Appendix Figure E2, the share of high-skilled workers in the model population rose from 22.5% in 1986 to 49.1% in 2017.

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

Table E1: Relevant dimensions for defining the model groups

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

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

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

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

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

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

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

Figure E1: Evolution of age group shares over time

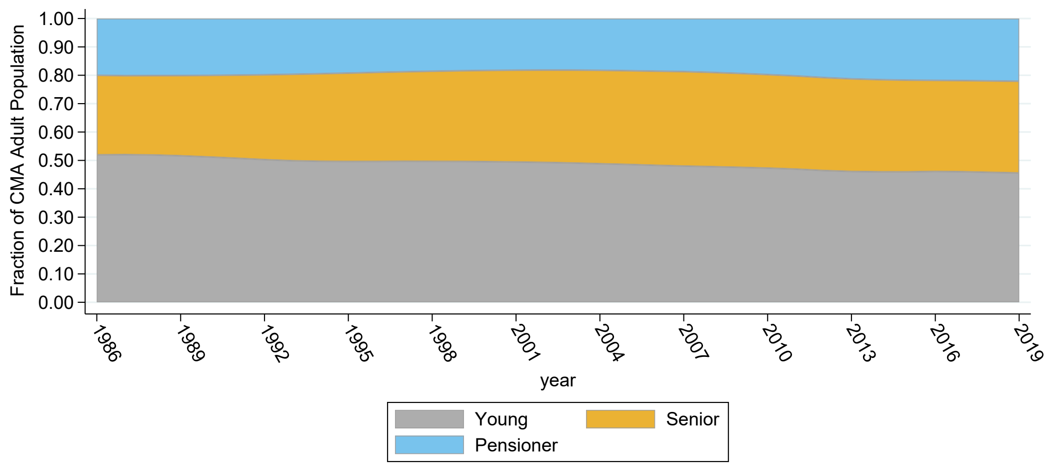

Notes: The figure plots the evolution of age group shares in the CMA adult population from 1986 to 2019.

Figure E2: Evolution of labor group shares over time

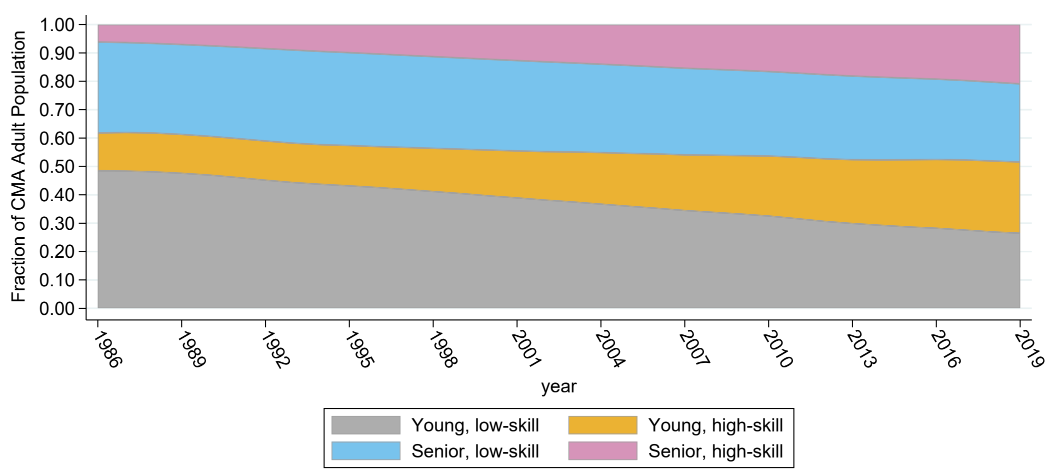

Notes: The figure shows the evolution of labor group shares in the CMA adult population from 1986 to 2019.

Figure E3: Evolution of family group shares over time

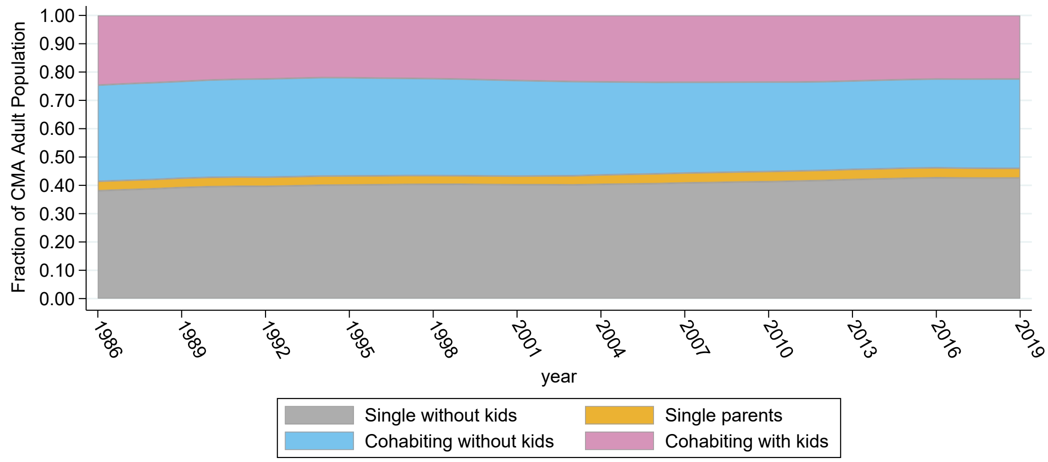

Notes: The figure shows the evolution of family group shares in the CMA adult population from 1986 to 2019.

E2 Parameter Estimation

Housing Expenditure Shares For the model calibration, we estimate group-specific housing expenditure shares, α_{of}. In our population panel, we observe aggregate household income net of taxation as well as residential location and floor space consumption. Using the universe of property transactions, we establish a local index of residential floor space prices in each year using the method of Ahlfeldt et al. (2021). To ensure that we are capturing location choices based on current market prices, we then restrict the sample to those that moved residence in the previous year. Lastly, we calibrate an annual yield to target an aggregate housing expenditure share of 30% for the total model population [8], and then estimate average housing expenditure shares for each group. The results are presented in Table 1.

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

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

Figure E4: Residential Amenity Comparisons

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

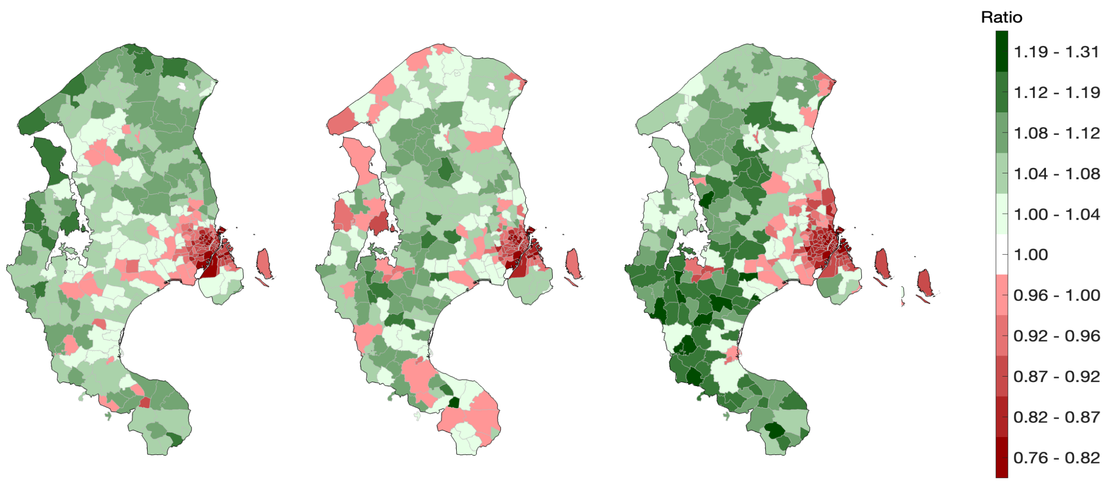

Notes: The figure maps the relative amenities between groups of different characteristics. The amenities for each age or family group are calculated with the weighted average across the worker groups with these characteristics. We show the ratio between average amenities of senior relative to young (Panel A), singles and cohabiting relative to cohabiting with kids (Panel B), and cohabiting relative to singles (Panel C).

F Demographic Counterfactuals

In this section, we describe how the demographic data for the Greater Tokyo Area[9] in 1980 and 2020 (presented in Table 3) was constructed. The main data sources for this table were the 1980 and 2020 Japanese censuses.[10]

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

F1 Japanese Census Data

Age Composition Table 2 of the 1980 Japanese census provides data on the total population by age (in single years) for each Japanese prefecture[11]. Using this data, we collect the number of people in the 18-44, 45-64 and 65+ age brackets and construct shares for these categories. Note that our dataset excludes everyone under 18 and people whose age is listed as unknown in the census[12].

Table 2-1-1 of the 2020 Japanese census provides the same data for all Japanese prefectures in 2020[13]. We use the same process to construct shares of the 18-44, 45-64 and 65+ age brackets in the total adult (18+) population in the Greater Tokyo Area.[14]

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

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

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

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

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

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

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

while all other individuals are classified as single (excluding people whose marital status is listed as "Unknown").[15]

Table 5 of 1980 Japanese census[16] provides marital status data for all 15+ individuals in each Japanese prefecture by 5 year age groups (15-19, 19-24, etc.), while Table 4[17] provides marital status data for all 15+ individuals across Japan by age (in single years). Using the first dataset, we collect marital status data for the 20-44, 45-64 and 65+ age groups in the Greater Tokyo Area. To include the 18-19 age group, we use the Japan-wide data to identify the fraction of singles and couples in the 15-19 age group that are 18 or 19 years old.[18]We multiply these fractions by the number of singles and couples in the 15-19 category in the Greater Tokyo Area to estimate the number of singles and couples aged 18-19 in the Greater Tokyo Area. We then aggregate across the age groups to obtain the total number (and share) of 18+ singles and couples in the Greater Tokyo Area in 1980.[19]

Table 4-1 of the 2020 Japanese census provides marital status data for all 15+ individuals in each Japanese prefecture by age (in single years)[20]. We use this dataset to identify 18+ singles and couples in the Greater Tokyo Area in 2020.[21]

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

The 1980 census broadly classifies households into 2 types - "ordinary households" and "quasi-households". An "ordinary household" is defined as a group of persons sharing living quarters (typically families) and living expenses as well as a person who lives by himself occupying a dwelling house, while all other households are classified as "quasi-households"[22].

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

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

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

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

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

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

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

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

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

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

Table F1: Ordinary Household Classification

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

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

Table 24 of the 1980 Japanese census[23] lists households, household members and household members under 18 years for each Japanese prefecture. We first remove the number of household members under 18 from the total number of household members for each household type to keep only the 18+ population in the Greater Tokyo Area. Using this dataset and our assumed number of parents per category as presented in Table F1, we calculate the number of parents in the Greater Tokyo Area by multiplying the number of households with the number of parents assumed per household for each "ordinary" household type. We then subtract the number of parents from the total number of 18+ household members for each household type to obtain the number of non-parents. Additionally, all individuals living in quasi-households are classified as non-parents, except for students living in school dormitories, who are excluded from the analysis.[24]We then aggregate over "ordinary" and "quasi-households" to obtain the total number of parents and non-

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

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

parents in the 18+ population of the Greater Tokyo Area in 1980.[25]

The 2020 census classified households into "private" and "institutional" households, which are broadly similar to the 1980 classifications of "ordinary" and "quasi-households" respectively.[26] The family structure classification of "private" households is identical to the 1980 classification of "ordinary" households, so we use the same assumptions listed in Table F1 to classify individuals as parents and non-parents in 2020.

Tables 9-1 and 9-3 of the 2020 Japanese census[27] list the number of households and household members (and household members under 18 years) respectively for each "private" household type in all Japanese prefectures. We follow the same approach as the 1980 data to filter out individuals under 18 years and classify the 18+ population into parents and non-parents in the Greater Tokyo Area. Note that, like 1980, all individuals living in institutional households are classified as non-parents, except for students living in school dormitories, who are excluded from the analysis.[28]

F2 Counterfactuals with Tokyo's Demographic Changes

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

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

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

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

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

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

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

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

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

Lastly, we simulate a combined counterfactual which implements all three of these demographic changes together. To implement the combined counterfactual, we sequentially impose the three demographic changes described above. Since each of the individual changes preserves the shares of the other two demographic dimensions, the order in which we implement the three changes has no effect on the resulting population distribution. [29]

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

Figure F1: Counterfactual Changes in the Workplace Employment

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

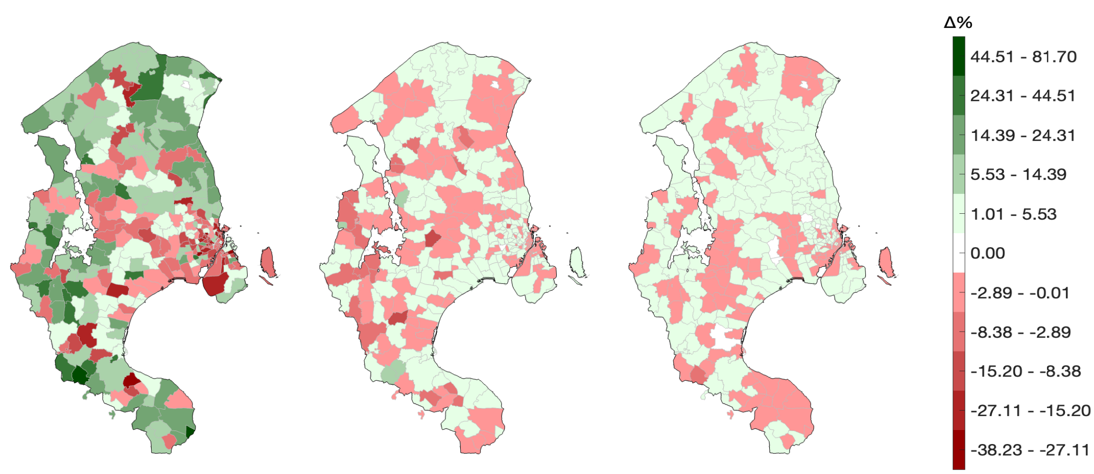

Notes: The figure presents counterfactual changes in workplace employment under three demographic scenarios: (a) aging, (b) fertility decline, and (c) increase in single households. Each panel shows the spatial variation in percentage changes in workplace employment across areas.

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

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

Figure F2: Counterfactual Changes in Residential Floor space Prices

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

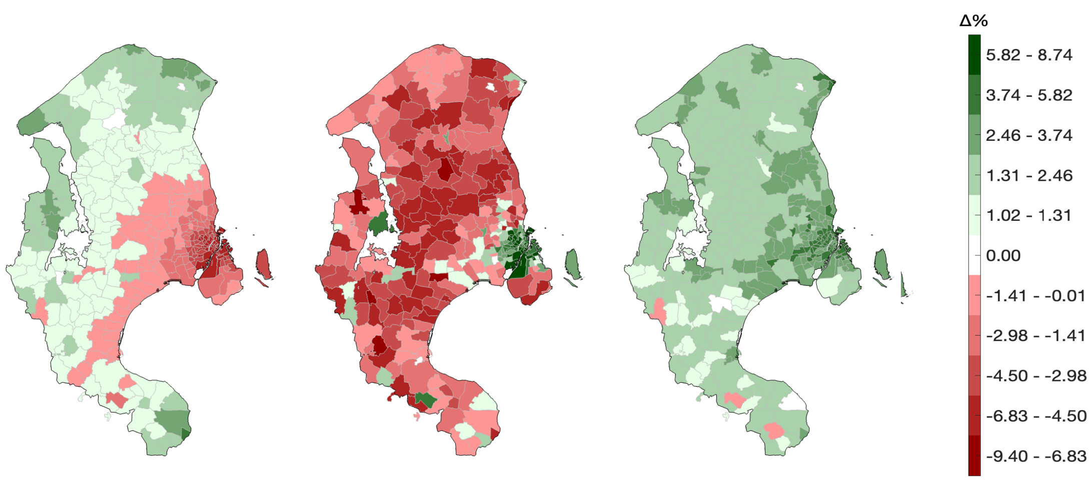

Notes: The figure presents counterfactual changes in residential floor space prices under three demographic scenarios: (a) aging, (b) fertility decline, and (c) increase in single households. Each panel shows spatial variation in percentage changes across areas.

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

Figure F3: Counterfactual Changes in Commercial Floor space Prices

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

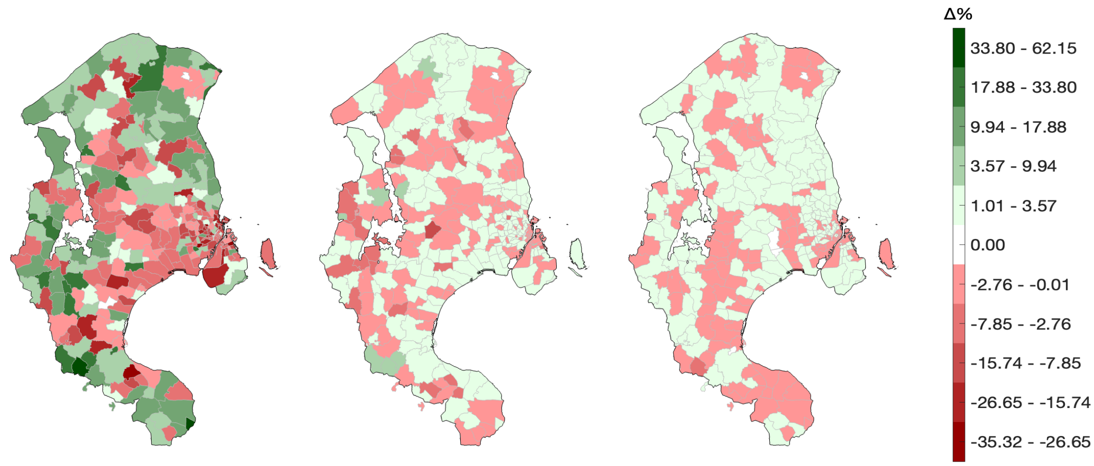

Notes: The figure presents counterfactual changes in commercial floor space prices under three demographic scenarios: (a) aging, (b) fertility decline, and (c) increase in single households. Each panel shows the spatial variation in percentage changes in commercial prices across areas.

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

Figure F4: Tokyo Counterfactual Changes

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

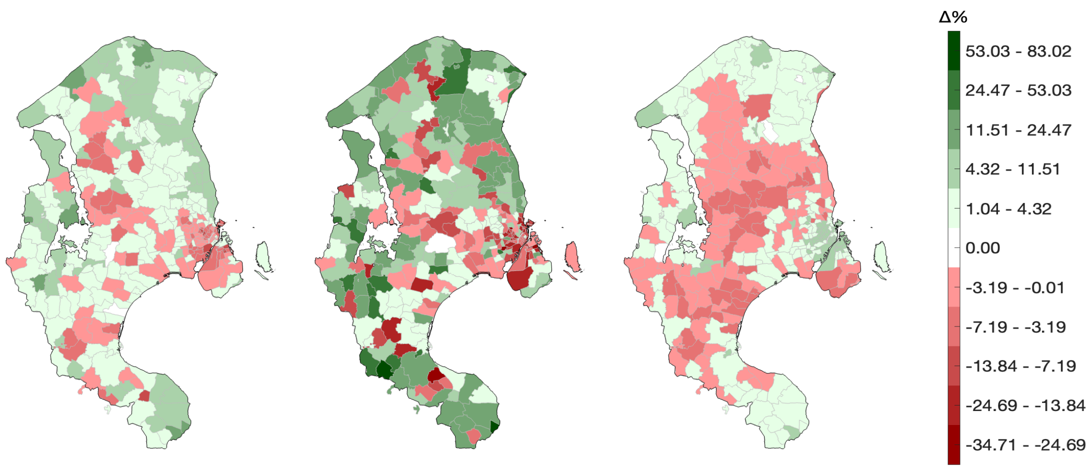

Notes: The figure presents counterfactual changes in Tokyo under three outcomes: (a) residential population, (b) workplace employment, and (c) residential floor space prices. Each panel displays the spatial distribution of percentage changes across areas.