UNIVERSITY OF COPENHAGEN
DEPARTMENT OF ECONOMICS

PhD thesis

Christian Langholz Carstensen

Three Essays on Housing Markets

Price Dispersion, Dynamic Location Choices and Family Investments

Advisor: Bertel Schjerning

Handed in: August 7, 2020

Contents

Acknowledgments

I remember quite clearly a summer evening around my final year of high school, sitting with my family around the dinner table. My parents were inquisitive about my intentions education-wise. I said it did not matter much exactly what I ended up working with, as long as it involved delving into some kind of analytical work. Something that involved focusing on getting to the bottom of a topic. It was not quite a helpful answer until my father clarified that I wanted to do research. This turned out to be true.

I am incredibly grateful of having been afforded the distinct luxury of academic life that is getting seriously into the weeds of a few interesting questions. It really was the hope of my teenage self, and I want to extend a heartfelt thanks to everyone who was part of it. First of all, I want to thank Peter Stephensen for taking a chance and hiring me for the PhD project at the Danish Research institute for Economic Analysis and Modeling (DREAM). It would not have happened without your engagement and I am very thankful for that. DREAM is a unique institution in the landscape of Danish policy advising and the years there have been very stimulating. The sometimes prolonged discussions we have had about agent-based modeling and other interesting but not necessarily related subjects have been rewarding and I really appreciate them. I also want to thank all the other fun, smart and quirky people at DREAM for the luncheoning, social gatherings and table football matches that spice up the every day.

Secondly, this thesis would not have been without the economic support from Realdania and the assistance of Curt Liliegreen at the Knowledge Centre for Housing Economics. I am honored that they chose to support this project and I sincerely hope that they will find the papers they have helped bring about worthwhile.

I must give a major round of thanks to Bertel Schjerning, my advisor at UCPH. I am taken aback by the great personal involvement you put into advising your PhD students, myself included. It has been a true privilege to be under your wings and I have learned a tremendous amount from these years (including the virtue of being stuck in a certain phase of psychological development when programming). It was also great fun to hang out with you and your family during the stay in D.C., and I am super glad that I got to know them. On that note, I also want to thank John Rust for inviting me to Georgetown and for the enormous hospitality. It was a delightful and intellectually stimulating time that I often look back on.

In addition to John and Bertel, I would like to thank Fedor Iskhakov and Maria Juul Hansen for taking me in on the URBAN paper (Chapter 2) and the work we have done together. The summerhouse boot camps have been exciting and hopefully the covid pandemic will not put an end to these adventures.

I am very appreciative of the funding provided by the department and I have been very happy to work there. The endeavor of writing a thesis would also have been much more dull if it was not all the great colleagues at UCPH, thanks to all of you for creating the great work environment that shapes the department.

Lastly, and most importantly, the greatest thanks goes to my girlfriend and son, who way too often have been putting up with an absentminded and stressed candidate. Your stamina is one for the books and I cannot wait for all the days, afternoons, weekends and holidays we are going to spend together.

Sammenfatning

Denne afhandling består af tre selvstændige kapitler, hvis omdrejningspunkt er husholdningers investeringsadfærd på boligmarkedet. Kapitlerne varierer i metode og emne, men fælles for dem er, at der er lagt vægt på de fordelingsmæssige konsekvenser af en given politik eller konkjunkturchok.

Kapitel 1 undersøger hvorledes et lokalt boligmarked kan opleve en differentieret udvikling i priser på tværs af boligsegmenter som respons på et konjunkturchok. Udgangspunktet for kapitlet er den observation, at priserne i Storkøbenhavn faldt mest på de i forvejen billigste boliger, da 00'ernes finanskrise satte ind. I den efterfølgende periode lukkedes gabet mellem prisindekserne for de billige og dyre segmenter, omend ikke fuldstændigt. Man kan desuden observere en tæt sammenhæng mellem boligers priser og likviditet målt i liggetid. At rationalisere denne udvikling på boligmarkedet kræver en ikke-stationær model, der både tillader søgefriktioner og samtidig rummer en høj grad af heterogenitet blandt husholdninger og boliger. Hertil opstilles en agent-baseret model, som netop er i stand til at opfylde disse krav. En kalibreret version af modellen reproducerer kvalitativt den udvikling i priser og likviditet, der efterfulgte finanskrisen i 2007 og kan således anvendes til at vurdere effekten af fremtidige konjunkturchok.

Kapitel 2 er skrevet sammen med Maria Juul Hansen, Fedor Iskhakov, John Rust og Bertel Schjerning. I dette kapitel præsenterer vi en dynamisk ligevægtsmodel for et boligmarked, der er fordelt over mange geografiske zoner med varierende offentlige goder. De geografiske zoner varierer derudover ved deres løn og ansættelsesmuligheder. En central innovation i modellen er, at husholdningerne foretager deres valg af geografisk bopæl i tandem med en beslutning om, hvor de vil søge ansættelse *uden* at disse behøver at dele placering. Husholdningerne vælger derved implicit i modellen, hvor langt de er villige til at pendle givet alternativerne. I tillæg fastholder modellen at flytninger er omkostningsfulde og har langsigtede konsekvenser. Husholdningerne er derfor modelleret som rationelle og fremadskuende i deres flytteadfærd. Vi estimerer modellen på mikrodata for Storkøbenhavn,

hvor vi lader kommuner udgøre den geografiske zoneinddeling, og simulerer to scenarier. Ét hvor udbuddet af boliger i Københavns centrum øges og ét hvor rejsetiden (eller, ækvivalent, omkostningen ved rejsetid) øges mellem alle zoner. Når udbuddet af boliger øges lokalt falder priserne overalt og urbaniseringen stiger. Idet rejsetiden øges falder priserne i udkantszonerne, husholdningerne flytter tættere på deres arbejdsplads og flere ender udenfor beskæftigelse.

Kapitel 3 sætter fokus på forældrekøbsordningen. Valget om at investere i en forældrekøbslejlighed indebærer en interessant økonomisk situation for forældrene, da de ved hjælp af denne både kan støtte deres børn finansielt, alt imens de selv har muligheden for at opnå et afkast. Afvejningen mellem forældrenes og børnenes interesser i investeringen er derfor sigende om forældrenes grad af økonomisk altruisme rettet mod børnene. Derudover subsidieres forældrekøbene via en række skattefordele. Regressioner på forældrenes investeringsvalg og nedslag i pris hvis børnene køber lejligheden af forældrene, viser at både altruistiske motiver og finansielle er i spil ved forældrekøb. Derudover opstilles en kollektiv model for investeringen i forældrekøb for at illustrere velfærdseffekterne af ordningen. Denne viser at skattefordelene øger uligheden blandt børn og har en tendens til at øge boligpriser.

Summary

This thesis consists of three independent chapters which all revolve around the investment decisions of households on the housing market. Although the chapters vary in method and subject, they all emphasize the distributional consequences of a policy or business cycle fluctuation.

In **Chapter 1**, I investigate how a local housing market may experience a differentiated development in prices across housing segments as a response to a financial contraction. The starting point of the chapter is the observation, that housing prices in Copenhagen fell the most among homes that were already in the cheapest segments before the onset of the financial crisis of 2007. The price index of the cheapest segments did somewhat close the gap to that of the most expensive segments in the following years, although they had not completely converged as of 2016. Furthermore, during the financial crisis and in the ensuing recovery, one can observe a close co-movement of housing prices and housing liquidity, measured in time-on-market. Rationalizing these trends requires a non-stationary model that allows for search frictions together with rich heterogeneity among households and homes. To this end, I create an agent-based model which is apt to meet these requirements. When calibrated to the housing market of Copenhagen, the model is able to replicate qualitatively the development in prices and liquidity following the financial crisis of 2007.

Chapter 2 is co-authored with Maria Juul Hansen, Fedor Iskhakov, John Rust and Bertel Schjerning. We present a dynamic equilibrium model for a housing market that is distributed across many geographical zones of varying amenity levels. Furthermore, the geographical zones also vary by their household specific wage and hiring opportunities. A central innovation of the model is that households make their residential location decisions in tandem with their job search decisions *without* assuming that job and residence must share geographical zone. Implicitly, households are thus choosing how far to commute. The model also underlines the costly nature of relocating which can have long term consequences. Households are therefore modeled as rational and forward looking when contemplating their work and residential

locations. We estimate the model on micro data for the Greater Copenhagen Area, letting municipalities constitute geographical zones, and simulate two counterfactual scenarios. One in which the supply of housing in Central Copenhagen is increased and one in which travel times (or, equivalently, the cost of travel time) are increased between all zones. As housing supply is increased locally, prices fall everywhere and urbanization is increased. When travel times increase, housing prices fall in peripheral municipalities, households commute less and more households end up without employment.

Chapter 3 considers the phenomenon of parental landlord arrangements in Denmark. A set of tax rules originally intended for self-employed workers benefit parents who purchase an apartment for the sake renting it out to their children. Such an arrangement places parents in an economically interesting trade-off, since it represents both an opportunity to reap financial gains while supporting their children. As such, the outcome is telling about the degree of economic altruism of parents. Regressions on investment choices and discounts if a child subsequently buys the apartment show that both motives direct parents' decision. In order to illustrate the welfare effects of the current tax rules, I construct a collective model of the decision process. It shows that the tax benefits do increase inequality in welfare among children and work to increase housing prices.

Chapter 1

Dynamics of Price Dispersion in the Housing Market

This paper presents a simulated model of price dynamics in an urban housing market encompassing important lines of heterogeneity among households and homes. Households are characterized by their income and mortgage position and homes by a measure of their quality. This enables the model to replicate cross-sectional effects of a housing cycle. These effects include the number of either defaulting or insolvent, and therefore immobile, households during a bust and capital gains during a boom. An important application of the model is that it is informative about the variation in prices across segments of the housing quality distribution when credit conditions change. The model is therefore applied to a scenario resembling the housing bust starting in 2007 in the Copenhagen Metro Area to explain the widening gap in prices of low and high quality homes.

This project was enabled by funding from Realdania and the Danish Research institute for Economic Analysis and Modeling which is gratefully acknowledged.

1 Introduction

The 2000s boom-bust cycle in housing markets across the U.S. and Europe has spurred a voluminous academic discussion on the nature of a housing cycle. Much attention has been directed towards the effects of credit standards and interest rates on aggregate housing prices, see Davis and Van Nieuwerburgh (2015), but less has been said about the nature of a cycle across segments within a housing market.

This paper contributes to the literature on housing cycles by focusing on the cross sectional changes in prices that may follow in the wake of a shock to credit conditions. Heterogeneous price changes across housing segments have been documented empirically by Landvoigt et al. (2015) and Piazzesi and Schneider (2016) for cities in the U.S. and by Mulalic et al. (2017) for the Copenhagen Metro Area (CMA henceforth). Mulalic et al. (2017) find that high quality apartments have appreciated more in price since the end of the 2007-2009 financial crisis. A very basic hedonic approach confirms this finding in Figure 1 of the next section. It shows an increase in price dispersion during the bust stemming from the fact that home prices in the lower segments of the quality distribution took the hardest hit. It also shows that price dispersion have steadily decreased since then as prices in low segments are rebounding faster than in the top segment.

The contribution of the following analysis is to rationalize this pattern of price dispersion in tandem with the spiking time to sell that was also observed during the crisis. Doing so requires a model that can be simulated over a reasonable amount of periods, encompassing fine-grained heterogeneity among both households and homes.

A dynamic model of the housing market with these features should very quickly be rendered intractable if decisions are based on rational expectations in a Markov perfect equilibrium. The dimensionality of the problem quickly explodes and, accordingly, papers concerned with housing liquidity have usually disregarded heterogeneity among homes, see e.g. the review of the housing search literature by Han and Strange (2015).

In so far one is preoccupied with both price dispersion and liquidity in a combined framework, a way forward is to relax the notion of rational expectations and replace it with adaptive learning. Such behavior has been found when surveying home buyers by Case et al. (2012) and investigated in lab experiments by Bao and Hommes (2015), Hommes (2011). Glaeser and Nathanson (2015) further argue that the typically observed autocorrelation in housing prices cannot be accounted for without the presence of extrapolating price setters.

I follow this idea in creating a simulated model of a local housing market which comprises a large number of households who trade homes with each other. There is only a fixed set of homes in the economy so that whenever a households desires to change its housing position, it must do so in trading with another household. Homes are characterized by a scalar measure of quality, while households are characterized by a stochastic income and their accumulated home equity. Trades are subject to search frictions (in the sense of directed search) that, together with adaptive price expectations, creates auto correlated prices. The model is calibrated to the CMA using micro data and does a quite reasonable job in matching the observed cycle in price dispersion and selling times.

Related literature

An important feature of the simulation is that households act and trade homes without being subject to an explicit market clearing condition. As we will see, this does not entail that the market cannot find an equilibrium. It only means that the equilibrium occurs after a period of adaptation when conditions change due to the learning process of agents. As such, the model may be termed an agent-based model, see e.g. the discussion in Sinitskaya and Tesfatsion (2015).

A similar approach has been taken by Geanakoplos et al. (2012) who study the role of leverage and diminished credit standards in the booming housing market of Washington DC. Based on that paper, Baptista et al. (2016) at the Bank of England developed a model for the effects of credit standards in the UK housing market. Also closely related to the agent-based literature on housing prices is Khandani et al. (2013) who study how refinancing, low interest rates and rising home prices in combination is likely to create systemic instability.

Another strand of literature which ties up closely with the present model is the assignment models of a housing market, developed in Määttänen and Terviö (2014) and Landvoigt et al. (2015). The basic idea of those papers is that imposing a market clearing within a specific time interval, say a year, entails an equilibrium set of prices that must ensure buyers being sorted perfectly into homes according to quality and buyers' wealth. Solving for such a price schedule allow them to understand how capital gains need not be uniform across housing segments. From section 5, it is argued that the current model has much the same structure as an assignment model, but in a dynamic setting with liquidity risk added.

The remaining paper consists of the following: Section 2 in which some important empirical facts about the recent boom-bust cycle in the CMA are presented to motivate the structure of the model. Section 3 describing the model it self and Section 5 which explaining aspects of the equilibrium. Section 4 which describes how the model is calibrated to the CMA and finally in Section 6 the model is applied to recent bust cycle in the CMA.

2 Motivating Facts

In this section, I present some empirical observations on the 2000s housing cycle in Copenhagen that motivates the structure of the model. The data applied stems mainly from full-population registers of housing transactions and buyer characteristics obtained from Statistics Denmark. A more detailed description of these is provided in Appendix A. The data reveals a strong connection between sales prices and the liquidity of homes, between the effective mortgage rate and prices, and that prices have not reacted uniformly across the segments of the housing market.

As mentioned in the introduction, a striking feature of the 2000s housing cycle in Copenhagen is that the impact of the down turn was much less uniform across the tiers of the housing market than the upturn. This is displayed in Figure 1, which shows that the bottom

Figure 1: Price development in each quality segment

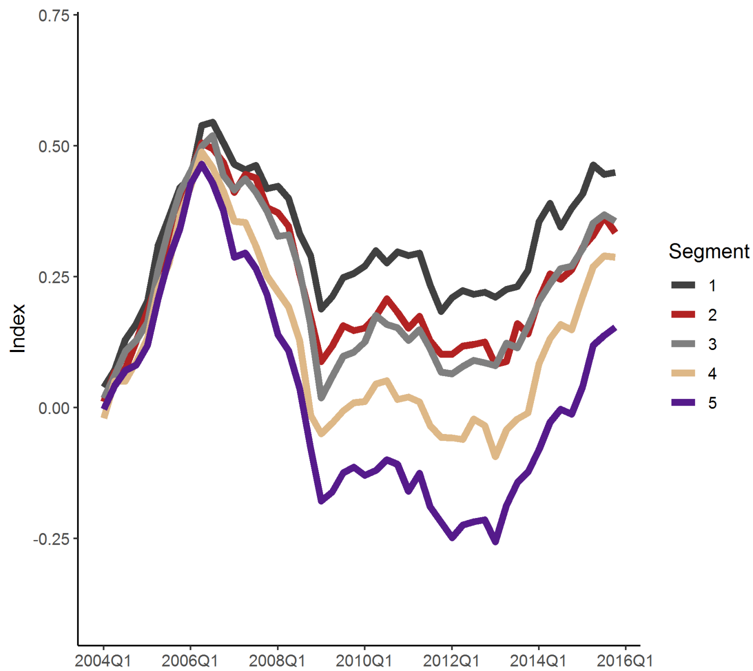

segment of the housing market both fell much more than the top segment and kept stagnant for a longer time. The fast growth in home prices seemed however to be uniform across segments in the boom years.

In order to construct Figure 1, the first step is to create a one-dimensional measure of quality for each home sold. In line with Landvoigt et al. (2015), one can use transaction prices in a base year to create a measure of quality given that the ranking of homes on average remains unchanged in the periods following the base year. The approach taken by Landvoigt et al. (2015) is to use a repeat sales model, however this requires more transactions than available for the CMA.

Instead, I apply a hedonic regression in the base year 2003 and use the obtained coefficients to impute the quality of each home that is sold in the resuming years. Based on this imputed measure of quality and contemporary sales prices, I then construct a price index for each market segment. Note how this approach assumes that changes in trading prices after the base year stem from changes in credit conditions and wealth but not from changing preferences.

Table 1: Percentiles of the housing quality distribution.

Percentiles				
20th	40th	60th	80th	100th
1,263,082	1,622,070	1,986,558	2,405,967	7,705,791
Fitted log-normal distribution				
$\hat{\mu}_q =$	14.387^{***}	$\hat{\sigma}^2_q =$	0.3788^{***}	

This assumption is likely not warranted for the neighborhoods that were intensively developed in the 2000-2010s, yet the conclusions remain if these areas are omitted. The hedonic regression in the base year has the following form

$$\log p_j = \delta_z \mathbb{W} + \delta_s sqm_j + \delta_h \mathbb{1}_{j=\text{ house}} + \delta_{1900} \mathbb{1}_{\leq 1900} + \delta_{1950} \mathbb{1}_{\text{1900-1950}} + \delta_{2000} \mathbb{1}_{\geq 2000} + \varepsilon_j. \quad (1)$$

The first term $\mathbb{W}$ contains a set of fine-grained neighborhood dummies (ie. the commuting zones described in Appendix A) to capture local amenities. The second term captures the effect home size above a minimum of 30 m^2. The third term captures whether housing is detached and the remaining three parameters are indicators of the home being constructed before 1900, between 1900 and 1950 and after 2010. Homes before 1900 are of historical value and often well situated while homes constructed between 1900-1950 were generally of higher standards than those constructed during the 60's and 70's building boom.

From the the predicted prices $\hat{p}_j$ of equation (1) we can construct a distribution of housing quality, from which the 20th, 40th, 60th, 80th and 100th percentile is obtained. These percentiles are presented in Table 1 together with the parameters of a fitted log-normal distribution.

Applying the estimated coefficients from the hedonic regression allow us to impute the quality of any home k sold subsequently to the base year by

$$\log q_k = \hat{\delta}_z \mathbb{W} + \hat{\delta}_s sqm_k + \hat{\delta}_h \mathbb{1}_{k=\text{ house}} + \hat{\delta}_{1900} \mathbb{1}_{\leq 1900} + \hat{\delta}_{1950} \mathbb{1}_{\text{1900-1950}} + \hat{\delta}_{2000} \mathbb{1}_{\geq 2000} \quad (2)$$

and then subsequently place any home k within a quality segment by the percentile cutoffs. The within segment development displayed in Figure 1 is finally obtained by running the following regression of time dummies on price relative to quality for the sales of each segment $\jmath$

$$\log(\frac{p_{kt}}{q_k}) = \delta^{\jmath}_{2004Q1}\mathbb{1}_{t=2004Q1} + \delta^{\jmath}_{2004Q2}\mathbb{1}_{t=2004Q2} + \cdots + e_{kt} \tag{3}$$

and then plotting the estimated $\delta^{\jmath}_t$ coefficients.

Figure 1 shows that the growth rate in prices before the bust were quite uniform (and hefty) across segments, yet prices fell more in the lower segments than in the upper relative to their price before the boom. Prices in the lower segment do nevertheless appear to be closing the gap in relative prices by growing faster than the top segment from around 2013.

An important component in the dramatic price development is the effective interest rate on newly issued mortgages depicted in Figure 2. The rise of the interest rate sets off quite precisely when housing prices are also starting to turn and it peaks when prices reach a low point in 2009. In addition, the period after 2012 is characterized by continuously falling interest rates and rising prices.

Figure 2: Average effective interest rate on newly issued mortgages

Notes: The displayed rate is a compound of actual interest payments, fees associated with obtaining a mortgage and contributions to mortgage institutions. Furthermore, the statistic compounds both the expenses of adjustable rate and fixed rate mortgages and all adjustment periods. Source: Central Bank of Denmark, DNRNURI.

The causality could of course run in the other direction; interest rates might be rising due to poor prospects on the housing market. In that regard, it should be noted that default rates were low and only modestly elevated throughout the financial crisis due to the recourse structure of Danish mortgages. Also, the drop in effective mortgage rate from 2009 must have been guided by monetary policy rather than market conditions, as it occurred exactly when the housing market was in its worst of shapes.

In light of the dramatic drop in prices from 2006 to 2009, it is interesting to consider the liquidity of housing markets in the same period. The website boliga.dk contains the vast majority of homes posted online for sale in Denmark by collecting all adds of real estate agents. The website also tracks previously posted prices for the same add and one can therefore obtain dataset of posted price adjustments by scraping this website.

Figure 3: Liquidity of apartments in the CMA

(a) Correction of posted prices

(b) Time on market

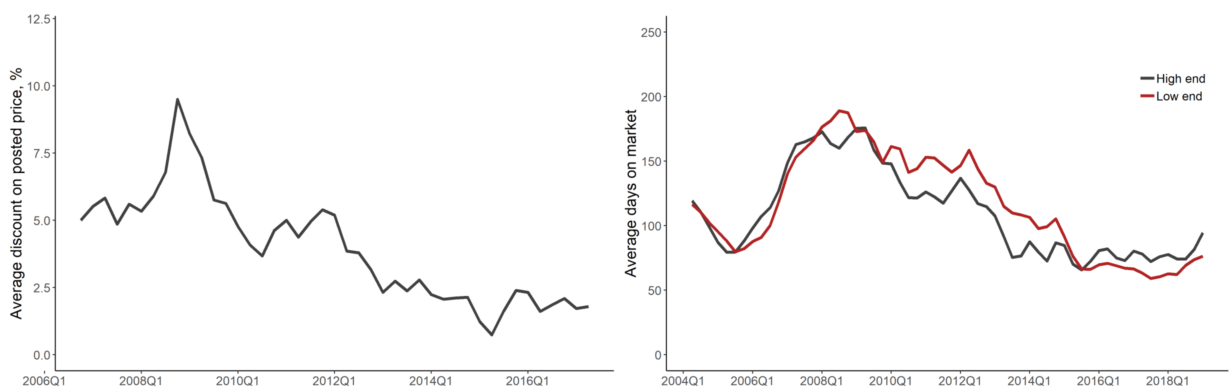

Notes (a): The figure shows the average difference between initial posted price and posted price at the time of removal of housing add. Source: Historical housing adds scraped from www.boliga.dk.

Notes (b): Data stems from zip code level average days on market before selling for apartments published by Finance Denmark, BM031. The high end of the market is defined by the set of zip codes with highest average quality which together span roughly 30% of all sales over the period. Likewise is the low end of the market defined as the zip codes that comprise the lowest average quality and taken together spans roughly 30% of all sales.

Panel (a) in Figure 3 shows the average reductions in posted prices since 2006. Notice how closely they follow the path of prices in Figure 1. Reductions spiked to 10% in 2009, where prices were plummeting; they came down to 5% in 2010 while prices were stabilizing and when prices took a minor dip in 2012, reductions went back up. Finally, as prices have been on the rise since 2014, reductions have gone steadily the other way. An important aside is that the the posted price of real estate is a binding floor in Denmark in that sellers using a real estate agent agree to sell to any buyer that bids above the posted price. This obvious sign of low liquidity in the market during the crisis can be complemented by the time to sell as depicted in Panel (b) of Figure 3. Time to sell also peaked during the bust, averaging 200 days, and was slightly lower in the high segments of the market during the recovery years. As noted above, this coincides with high segment prices rebounding faster than low segment prices.

3 Model

3.1 Households

The model features an overlapping set of N^h finitely lived households who experience a life cycle of earnings and housing purchases. The households are indexed by i and are at any point t characterized by their age age_t, income y_t and wealth w_t. The households reside in what is thought of as a city that comprises N^d homes, one for each household, such that $N^h = N^d$. Homes are indexed by h and characterized by an index of quality, q. The quality of a home should be thought of as a composite of the home's size and a local amenity level that is fixed throughout the simulation horizon. Households can occupy only one home at a time for residential purposes, and it is assumed that homes are not malleable by their owners. Homes do not depreciate and all homes in the economy are assumed to be freeholds so that renting is not an option. Households can finance their

Households face an age dependent probability of death, $d(age_t)$, which is given by a Gompertz hazard function. Whenever a household dies, it puts its home up for sale and exits the economy after it has traded off the home. In order to keep the number of households constant throughout the simulation, a newborn homeless household enters the economy as soon as the deceased household submits its home to the market. This household starts searching for a home to purchase immediately after birth.[1]

Apart from housing consumption, households have available a composite consumption good, c, whose price is numerated to 1. As noted above, households require a home to live in meaning that all substitution happens between home quality and consumption at current home prices. Household preferences over c and q are given by the Cobb-Douglas utility function; this is standard and motivated by the fact that the expenditure share on housing has remained fairly constant over time.[2]

$$u(q, c; h) = (q_h)^{\alpha}(c)^{1-\alpha}. \tag{4}$$

The life cycle income of households follows the structure provided in Heathcote et al. (2010). There is a deterministic age profile in income given by $f(age)$ and a permanent-transitory stochastic process $\tilde{y}$. Household income y_t thus follows the process

$$\log y_t = f(age_t) + \tilde{y}_t \tag{5}$$

$$\tilde{y}_t = z_t + \epsilon_t \tag{6}$$

$$z_t = \rho z_{t-1} + \eta_t \tag{7}$$

The transitory shocks are log-normally distributed with $\epsilon \sim \log \mathcal{N}(0, \sigma_{\epsilon}^2)$ and permanent shocks with $\eta \sim \log \mathcal{N}(0, \sigma_{\eta}^2)$. At age^0, when a household is 'born' and enters the housing market, income is given by $f(age^0) + \epsilon_t + \eta_t$. In order to calibrate the model, this process is defined on an annual basis. Since the time unit of the simulation is months, the income of a household is updated once every 12 periods and kept constant meanwhile; note that this update is not synchronized across households. Note also that although the income measure in (5) is annual earnings, for the remainder of the paper y_t designates that measure divided

[1]Furthermore, the initial age distribution of households is calibrated against the death probabilities to ensure that there are no demographic waves through out the simulation.

[2]See e.g. Berger et al. (2017), Landvoigt et al. (2015) or Favilukis et al. (2017).

by 12 to get monthly income.

The point of modeling household income as stochastic is to give owners a motive to reoptimize from time to time by trading with one another. This is implemented by letting owners face a constant probability of becoming attentive to current market conditions, denoted λ^a. If market conditions seem favorable enough, in a way that is described below, the attentive household will put her home up for sale and start searching for a new home to purchase. As is often the case with models of market micro structure, a certain mass of 'noise traders' are needed in addition to those who act voluntarily, so as to not let the market come grinding to a halt. In the present case, one may end up in a situation in which all agents await each other to make a trade that reveals current prices. The traders who ensures that this does not happen are owners hit with a forced moving shock, occurring with λ^f, compelling them to put their home for sale irrespectively of their current income and home quality. Such moves are usually motivated by events in family relations like divorce or change in number of children living at home.[3]

As households enter the economy at age age^0, they carry an initial wealth endowment, w_i^0, which they use to meet the down payment requirement on their first home purchase. Initial income and wealth is assumed to follow the joint distribution F_{yw}. The dynamic consumption-savings problem of households is not explicitly formulated in this model, since incorporating rational expectations in a simulation such as this is computationally infeasible. It is instead assumed that household wealth is fully comprised by home equity, which is increased by paying installments as well as through capital gains when prices rise.[4]

Home purchases in the model are financed with a fixed rate mortgage that carries an option to defer installment payments until next period and only pay the interest rate expense. Households have two motives for this; one is that income may fall so much that residual income

[3]The attention shock is realized before the forced shock. If both obtain and market conditions are not favorable, the household will still have to move.

[4]There is a special case in which households do use another savings vehicle than home equity. If a new born household cannot meet the down payment requirement on even the cheapest home on the market using its initial wealth endowment, it will save up αy_t in cash for each period that it cannot afford any home.

after paying installments leaves them with consumption below a critical value c^{min}. The other motive is based on the unmodeled life cycle dynamic of savings. There is empirically a clear age gradient in interest-only payments as households have little interest in accumulating more housing wealth after a certain point in life. Younger households are also less inclined to pay installments due to intertemporal consumption smoothing. Thus, in order to keep home equity at an empirically relevant level, the function $g(age_t)$ defines the purely age-dependent probability of interest-only payments. The decision to pay only interest on a mortgage, $io = 1$, is therefore

$$P(io_t = 1 | age_t, y_t) = \begin{cases} g(age_t) & \text{if } y_t - hc_t > c^{min} \\ 1 & \text{else,} \end{cases} \tag{8}$$

where hc_t is the combined installment and interest payments on the current mortgage of a household. Note that there is no refinancing option in the model, so the effect of cashing out on equity must be captured through $g(age_t)$.

The level of home equity becomes important for the dynamics of the housing market both through the demand and supply side. In times of falling prices, home equity vanishes rendering buyers less able to fulfill the current loan-to-value requirement. On the supply side, sellers are unable to sell below the principal on their outstanding mortgage debt due to the full recourse structure of Danish mortgage debt. Low equity households therefore risk being unable to go on the market during a bust, adding to a contraction in supply.

Finally, it should be stressed that the problem of selling and buying is divided into two separate subproblems. This is in a sense a behavioral heuristic as the selling and buying process is deep down a joint problem for the household, see Moen et al. (2015), but that is typically simplified away for the sake of tractability in the search literature. When searching for a home to buy, the objective of a household is to purchase the home that yields the highest possible utility at current market conditions. When selling a home, the objective is to earn as much as possible while minimizing costly time spend in the process. This behavior is developed further below.

Seller behavior

As mentioned, the housing market is modeled as a posted price market. As such, a seller must figure out what her posted price should be when putting her home on the market. Here, we assume that sellers estimate the going rate for their homes by considering what other homes in the same quality range have been trading at within a recent number of periods. Using information in recent local sales to set the asking price of a house appears to be a reasonable behavioral assumption based on the degree of momentum in housing prices as well as guidelines for home evaluation posted websites for Danish mortgage lenders [5] as well as Glaeser and Nathanson (2015) and Case et al. (2012).

When buyer i, who is selling a home of quality q in t, estimates the home's current market price $\hat{p}(q)$, she does so utilizing a LOESS regression on all sales that have taken place during the past ς periods. This set is denoted $\mathcal{H}_t$. The LOESS regression is a weighted local regression that weighs observations by their nearness to the single point at which one would like predict the outcome variable. In the current setting, a seller wants to predict the market price at q. The fact that the LOESS weighs more distant observations less is an attractive feature for the seller as she realizes that the price may not be well approximated by a linear function of quality when moving sufficiently far away from q. The objective of the LOESS regression centered at some point q_0 is

$$\min_{\beta} \sum_{q \in \mathcal{H}_t} w(q)(p - \hat{p})^2 \tag{9}$$

$$\hat{p} = \mathbf{q}\beta, \quad \mathbf{q} = [1 \quad q]$$

$$w(q) = \begin{cases} \left(1 - (\frac{|q-q_0|}{d})^3\right)^3 & \frac{|q-q_0|}{d} < 1 \\ 0 & \text{else,} \end{cases}$$

[5] See e.g. https://www.rd.dk/da-dk/privat/Kundeservice/Spoergsmaal-og-svar/grundlaeggende-om-realkredit /Pages/hvordanVurderesEjendom.aspx

The bandwidth parameter d determines the degree to which distant observations are included in the prediction.

It is computationally heavy to estimate the LOESS every time a seller needs to predict the price of a home. To ease the computation, a set of centroid LOESS coefficients are estimated at evenly distributed points, $\{q_1, q_2, \ldots, q_G\}$ in the full domain of home quality. The price that a seller will have estimated is therefore the interpolated price between the two nearest centroids. The interpolated price for a home of quality q, which fulfills that $q_j \leq q < q_{j+1}$, will thus be

$$\hat{p}(q) = w_0 \mathbf{q}_j \beta_j + w_1 \mathbf{q}_{j+1} \beta_{j+1}, \quad w_1 = \frac{q - q_j}{q_{j+1} - q_j}, \quad w_0 = 1 - w_1 \tag{10}$$

In putting their home up for sale, sellers are willing to make a trade off between selling the home quickly and earning a markup, as is standard in housing search models with posted prices. This trade-off is implemented by assuming a fixed period-wise cost of having a home on the market together with the option of setting an initial markup over the estimated market price for their homes. This really makes for a dynamic programming problem as described in Merlo et al. (2015); yet the fully fledged dynamic optimization cannot be solved practically in a model such as this due to its high dimensional state space and lack of explicit equilibrium conditions. The solution is therefore that sellers must learn from observing each others' past behavior and outcomes so as to adapt their behavior to current market conditions. The learning mechanism is inspired by the fact that the relation between initial prices, final prices and time on market is public information in Denmark and available on free websites.[6] Before delving into the learning process, assume that a seller puts her home on the market in t_0 and has decided on a markup γ^u. Then the posted price will be

$$p_{t_0} = (1 + \gamma^u) \hat{p}_{t_0}(q). \tag{11}$$

[6] eg. www.boliga.dk

For each period the seller's home is not sold, it downgrades the asking price by $e^{-\gamma^d(t-t_0)}$ to attract more buyers, although not below the principal of the mortgage, m_t

$$p_t = \begin{cases} p_{t-1}e^{-\gamma^d(t-t_0)} & p_{t-1}e^{-\gamma^d(t-t_0)} > m_t \\ p_{t-1} & \text{else.} \end{cases} \tag{12}$$

This pricing behavior reflects the notion that sellers become increasingly desperate as time goes by without selling due to unmodeled practical challenges of having a home for sale that causes them to lower the reservation at increasing rates. The shape of this reservation price of sellers is chosen to roughly represent that of Merlo et al. (2015). They show, in a rational forward looking setting, that the reservation price of sellers initially does not fall very much for every unsold period, yet eventually starts moving quickly. Note that this behavior of sellers allow prices to respond to changing market conditions endogenously; if the price sensitivity of buyers fall due to increased incomes, change in credit conditions or lower utility of outside option then buyers will accept higher markups which will feed into the expected price of new homes for sale. If buyer conditions change for the worse, more homes will go unsold which increases selling time, entailing lower asking prices.

I assume that the initial markup is an object of social learning and adaption, but that the downgrading factor γ^d is fixed and common for all sellers. The choice of markup is discretized on a grid γ^u indexed by $l = 1, \ldots, K$. Each seller i in t observes all trades that has taken place since $t - \varsigma$. Denote this set of sales by $\mathcal{H}_t$. For each sale j, they observe the accompanying γ^u_j, the initial market price, the final price and how long it took. Assuming that households have homogeneous preferences for selling, we can construct a measure of fitness for a selling strategy, γ^u_l that depends on these observed variables. The constant, monetary cost of having a home on the market is denoted c^s. The price $\hat{p}_t(q_j)$ denotes the estimated market price for j in its initial posting period and T_j is the number of periods it took to sell j. By these metrics, we can formulate the fitness of the markup strategy that lay behind sale j by

$$\nu_j = p_{jt} - \hat{p}_t(q_j) - T_j c^s. \tag{13}$$

The first part, $p_{jt} - \hat{p}_t(q_j)$, is the profit earned on the sale and the second part the costliness of its selling time. Now, given that sellers observe all recent trades, each chooses to imitate the markup strategy γ^u behind one of the observed sales with a probability that is proportional to its fitness. Taking the exponential of (13) ensures that the measure remains positive, and the resulting logit probability of imitating the strategy of sale j is given by

$$P[imitate\ j] = \frac{\exp(\beta\nu_j)}{\sum_{l \in \mathcal{H}_t} \exp(\beta\nu_l)}. \tag{14}$$

The parameter β is the so-called greediness of the learning mechanism; if the greediness goes to infinity, the decision maker will put probability one on choosing the alternative with highest value.[7] It is also necessary for the learning dynamics that sellers experiment at random; otherwise, actions that have turned beneficial due to changing circumstances may go undiscovered if they were not chosen in the past. Hence, with a probability ω, instead of imitation, a seller choses between markups with a uniform probability. Versions of this learning mechanism has been shown to lead to stability and approximately rational behavior in a range of settings, see Arifovic and Karaivanov (2010), Anufriev et al. (2016), Salle et al. (2017), Gintis (2007).

Buyer search

The housing market works as a posted price market in which sellers set a take it or leave it offer for a home that buyers may choose to accept. When searching through the market, buyers act in the spirit of directed search in the sense that their encounters with sellers are not completely random but directed by prices. This is possible because buyers have extended

[7]Alternatively, one can think of it as the inverse scaling parameter of the random part of utility in discrete choice models.

knowledge about the stock of homes for sale. Due to the costly nature of partaking in a bidding round for a home, we assume that each buyer is only able to submit offers for at most J homes in every period, even if J is much lower than the number of homes that they have observed for sale. Importantly, we also assume that if any one seller of the J homes accept the offer submitted, then the buyer will be happy to purchase that home without waiting for the response of the other sellers. There are two aspects to this behavior; on the one hand it represents a high degree of risk aversion of buyers. Since they do not know whether any of the other trades they submitted bids for will come through at the time of seller's announcement, they choose to go through with the trade so as to not end up empty handed at the end of the period. Another interpretation is that buyers are satisficing in the sense of Simon (1956); ie. any outcome that is close enough to the first best alternative is acceptable to buyers. This is also a motivation for keeping J small.

Buyer search is implemented by letting each moving household i draw a sample $\mathcal{S}_t$ of houses for sale on the housing market at time t. For now, I let $\mathcal{S}_t$ comprise all homes for sale at t. Buyer i thus traverses through all homes in $\mathcal{S}_t$, calculating their implied utility and ranking them in order to pick the J most preferable ones to bid for.

In order to calculate the utility of each home on the market, buyers must first establish their budget constraint which includes home equity. A complication that arises in the model is that since households act both as buyers and sellers simultaneously, some households will sell their old home before they get to buy a new home and others vice versa. A household that buys before it sells cannot know for certain how much it will cash out from the trade. It must rely on its expected wealth

$$\hat{w}_t = \hat{p}_t(h) - m_t \tag{15}$$

when searching through the housing market. The expected price $\hat{p}_t(h)$ is obtained by equation (10) the same way as when selling. That households are able to buy before selling comes down to assuming either that final contracts are settled after the buyer has traded or that the bank is willing to extend a loan of size $\hat{w}_t$ at no costs to help the trade go through. For those households who do end up selling before buying, $\hat{w}_t$ is replaced with the actual cash out from selling as soon as the sale comes through. As will be clear in Section 3.2, the market clearing process implies that if household i's home is sold at time t (before it has succeeded in buying a new home), then i cannot update it's bids for homes until period $t+1$. It will not be able to submit new offers until $t+1$, even if they may have changed their willingness to pay through the wealth effect of retrieving a cash out that was different from the expected. In case i sells in period t, but succeeded in buying in $s < t$, it must add $\hat{w_{is}} - w_t$ to the principal on the mortgage which was created in s, thereby also correcting the monthly due payments.

Buyers are subject to the budget constraint that post tax income must cover housing costs in the form of mortgage payments, hc, and consumption of the composite good in every period

$$(1-\tau)y_t = c_t + hc_t, \tag{16}$$

where τ is the tax rate.

In order to determine a buyer's reservation price for each home, it is assumed that there is an alternative housing market outside the city with an infinite supply of homes at any q that trades at an exogenous price schedule. This can be interpreted as construction costs in the sense that there is a common price for homes of a minimum size. The price of larger homes is scaled up by a fixed marginal price pr square meter. By assumption, the price of quality is 1 in the outside option. Hence, quality and price of a home of size s is given by $q(s) = \delta_0^o + \delta_s^o s$, which is chosen optimally by a household given its perceived equity $\hat{w}_t$ (see equation (15)) and income. Note that the minimal house size is normalized to 0. By the preferences in (4), the optimal size is given by

$$s^* = \frac{(\frac{y_t}{\xi(r_t)} + \hat{w}_t)\alpha - \delta_0^o}{\delta_s^o}. \tag{17}$$

Quality chosen in the outside option is therefore $q(s^*) = \delta_0^o + \delta_s^o s^*$, which, since $p^o = q$, results in utility

$$u_t^o = \natural u(y_t - \xi(r_t)(p^o(s^*) - \hat{w}_t),\ q(s^*)). \tag{18}$$

The scaling factor $\natural < 1$ represents the notion that house quality in outside option (mostly characterized by size) cannot fully substitute for city amenities.

Inspired by the intuition behind Rosen's bid function, Rosen (1974), and the Iterative Bidding Algorithm of Kuminoff and Jarrah (2010), define the reservation price of a buyer (in terms of monthly expenditure) for home h of quality q as the variation in consumption, $\tilde{c}$, that leaves her indifferent when going from the outside option to h. The reservation price thus follows from assuming equality between u_t^o and the utility of h and inverting the utility function

$$\tilde{c}_t(q_h) = \left(u_t^o q_h^{-\alpha}\right)^{\frac{1}{1-\alpha}}. \tag{19}$$

Using the budget constraint in (16) and (19), we can rearrange to get a buyer's reservation expenditure for h as

$$hc_t^r(h) = y_t - \tilde{c}(q_h). \tag{20}$$

Finally, taking this reservation expenditure together with the cost of borrowing and equity will pin down the reservation price. The interest rate on mortgages, r_t, is common to all households. Given r_t and the number of terms to maturity (30 years), the monthly payments of a fixed-rate mortgage can be calculated by an annuity factor on the selling price, which is denoted $\xi(r)$. Therefore, in making the down payment w_t, monthly housing costs of a house priced at p are given by

$$hc_t(h, w_t, \xi_t) = \xi(r_t)(p_t(h) - w_t). \tag{21}$$

Substituting $hc_t^r(h)$ of equation (20) into (21) and rearranging thus returns i's implied reservation price for h as

$$p_t^r p^r(h; \hat{w}, \xi) = \frac{hc_t^r(h)}{\xi(r_t)} + \hat{w}_t. \tag{22}$$

An important note is that the Danish mortgage market is somewhat different from, say, the

U.S. with respect to the defaulting options of households. In addition to securing loans with the underlying property as collateral, households are personally liable for their mortgage debt; i.e. there is no walking away from homes. As a result, there was only a modest increase in defaults following the recent financial crisis; around 0.5% of the housing stock went into foreclosure annually in the aftermath of the crisis. Those who did in fact default were accordingly in a financially dire shape at the time of default, Haldrup et al. (2015).

Whenever a collateralized home is traded, the mortgage will usually have to be redeemed by the borrower. It is legal, yet completely at the discretion of the lender to allow, that the new owner takes on the existing mortgage of the previous owner. This may for example be desired in case of a divorce, where both parties are liable for the mortgage and own a share of the home. In case of financially strained owners, lenders are not likely to allow further leveraging. The full recourse structure of Danish mortgages creates a high degree of certainty for lenders allowing them to profitably operate with a high collateralization rate. The downside is that it also bars under water households from selling, which deteriorates market liquidity in case there is a coordinated drop in income and housing prices.

Due to the high costs of defaulting, it is therefore assumed that households will not sell, despite drawing a moving shock, if the expected price of their home is below their outstanding mortgage debt. Upon receiving a moving shock, the household thus needs to check whether it is under water before going on the market. This is done by forming an expectation of the current market price for its home h, $\hat{p}(h)$, through the same method of interpolation between price centroids as sellers do in (10). In so far as the principal on the household's mortgage, m_t, does not exceed the expected selling price, $\hat{p}_t(h) > m_t$, the household is willing to go on the market.

Due to the finite number of agents in the economy, some houses may become overly attractive due to price experimentation resulting. This causes a majority of buyers to submit bids for the same set houses, which dries up liquidity in the market. To smoothen the demand for any single home, an idiosyncratic taste shock is added to each home household i inspects on the market. The quality perceived by i is therefore the sum of a common quality and a subjective term

$$q_{ht} = q_h + \varepsilon_t, \;\; \varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2). \tag{23}$$

As noted above, the objective of buyer i is to find the J most preferable homes in the choice set $\mathcal{S}_t$. The jth of these homes is thus the one yielding the highest utility in $\mathcal{S}_t$, excluding the ones better than j and obeying the associated constraints

$$u(q_h, c; h_j^*) \geq u(q_h, c; h), \quad \forall h \in \mathcal{S}_t \setminus \{h_1^*, \ldots, h_{j-1}^*\} \tag{24}$$

$$s.t.$$

$$c = y_t - hc(h, w_t, \xi_t) \tag{25}$$

$$c \geq c^{min} \tag{26}$$

$$p_t^r(h) \geq p_t(h) \tag{27}$$

$$\frac{w_t}{p_t(h)} \geq \theta^{ltv} \tag{28}$$

Equation (25) yields consumption as a function of housing expenditure, equation (27) is the requirement that asking prices are below reservation prices, (28) is the loan-to-value collateral constraint set by banking sector, and (26) is the requirement that composite consumption does not fall below a critical value. The minimal consumption constraint is not as strictly defined as the collateral constraint by regulators in Denmark since it is at the discretion of lenders to define a relevant minimal level. However, the minimal consumption level is restricted by the limits formulated by The Financial Supervisory Agency for classifying borrowers as high, normal or low quality in the balance sheets of lenders.[8]. Absent official data on lender policy, it is assumed that lenders will require that borrowers pass the FSA limit of normal quality at the time of granting the loan.

After collecting the set of homes for sale that the buyer finds satisficing, it submits a bid for each of them corresponding to the posted asking price. For reasons described below, there

[8]The legal guidelines are available in Danish in The Danish Financial Supervisory Agency (2018)

is a risk that the buyer will not get the opportunity to purchase any of the homes bid for and in that case, i repeats the process in the following period until success. The same applies if no homes on the market are affordable or if none of them are above i's reservation price.

Voluntary search

As noted above, household face the chance of receiving an attention shock to the market. The role of attentive movers is to make the volume of trades responsive to changing market conditions. When owners receive an attention shock, they follow the same process as outlined above with two modifications. Instead of browsing through all homes for sale, they test if current market conditions may provide them with a better alternative than their current home. As a heuristic for testing market conditions, they traverse through each of the commonly known centroids of the LOESS regression and calculate if buying a home of the associated price and quality yields a higher expected utility than the utility experienced in the currently owned home. If this is the case for any node, they decide to move. Denote the set of LOESS centroids at t by $\mathcal{L}_t$, and note that hc_t are mortgage costs including installments.

$$u(q_h, y_t - \xi(r_t)(p(h) - \hat{w}_t)) > u(q, y_t - hc_t) \quad \exists h \in \mathcal{L}_t \tag{29}$$

Obviously, this process abstracts away from the financial costs of moving which households empirically do face. The reason for this abstraction is mainly computational. A sufficient number of trades must take place every period in order for the market to keep going. For a computationally practical number of agents in the simulation, the propensity to move must therefore be higher at every age than the empirical counterpart. Yet it is not the objective of the model to match life-cycle moving patterns pr se, only to match the price distribution and allocation resulting from market conditions. In that sense, each agent represents more a type or set of households than a specific individual.

3.2 Market clearing

After all buyers have submitted their offers, sellers take turns in random order to accept incoming offers. If a seller has multiple buyers submitting offers on her home (which are all equal to the asking price), the home is allocated to the buyer with highest reservation price. When the trade is done, the buyer who won the bidding round retracts any bids submitted for other homes on the market. Then, the next seller in line accepts an offer, etc. As a motivation for allocating homes to the bidders with the highest reservation price, we note that an efficient sorting equilibrium implies that each home is allocated exactly to the household that is willing to pay the most for it, Kuminoff et al. (2013). Therefore, buyers will likely be better off with this allocation procedure compared to, say, a random allocation or allocating homes to the wealthiest buyer.

In Section 2, we saw that there were significant fluctuations in the liquidity of the housing market following the recent financial crisis. In order to allow for such fluctuations in liquidity, a search friction is build into the market mechanism that may keep it from clearing perfectly. As noted in Section 3.1, buyers submit only a limited number of offers for homes, meaning that favorably priced homes will attract more buyers at the expense of those with a high markup. The result is a time consuming selling process where the probability of receiving bids for one's home depends on the added markup as well as market tightness. A negative shock to credit conditions, as an example, will therefore induce more buyers to opt for the same small subset of cheap homes on the market, resulting in prolonged time-on-market and falling asking prices for homes outside the subset.

4 Calibration

The model is parameterized to be roughly consistent with the Copenhagen housing market in 2006. Incomes and prices are deflated by the consumer price index of Statistics Denmark with base 2011. All parameters are collected in Table 2 below. Part of the calibration relies on feeding into the model a distribution of housing qualities that match sales prices in 2004-2006 together with a set of agents whose income profile is estimated on micro data for the Copenhagen Metro Area in 2005-2011. This micro data stems from full population registers of Statistics Denmark. In particular, the registers BEF (containing information about residence, age and family status), INDH (income data) and REAL (mortgage data since 2012) have been combined.

Income process

The income process outlined in equations (5)-(7) is estimated using the strategy laid forth in Heathcote et al. (2010). Because homes are bought by individuals as well as couples, the income process assigned to households must be representative of both. I therefore consider couples as a single entity, summing their annual income and defining household age as the average of both partners. Annual income[9] is defined as the sum of salary, income from self-employment, capital income, public transfers and pension. I discard observations of households earning less than 100,000 DKK a year before taxes as well as more than 2 million DKK pr member. Discarded is also households living less than 4 years in CMA between 2000-2013. The resulting average life-cycle income profile is presented in Appendix B Figure 9.

Annual tax payments are also available in the micro register INDH, which is used to get the average tax rate. The average of tax payments over total income in the sample amounts to 0.36, which then defines τ.

Interest-only payments

The function $g(age)$ that determines the probability of switching to interest-only payments is estimated from the REAL register. The register comprises the stock of all mortgages in 2012 together with information on LTV, interest rate, person identifier, date of issuing and dates

[9] Using the variable PERINDKIALT of register INDH.

Table 2: Parameterization of the model

Parameter		**Value**
	Population	
Number of households	N^h	30,000
Quality of homes	q	$\sim \log \mathcal{N}(14.392, 0.378)$
Wealth to income ratio at birth	$\frac{y_i^0}{w_i^0}$	$\sim \mathcal{U}(3, 5)$
Variance of transitory shocks	σ_ϵ^2	0.00904
Variance of permanent shocks	σ_η^2	0.02208
Autocorrelation in permanent income	ρ	0.95
Tax rate	τ	0.36
Age of entrants to housing market	age^0	27
	Buyers	
Preference for housing quality	α	0.3
Minimal consumption	c^{min}	9,000
Moving shock probability	λ^f	0.015
Attention shock probability	λ^a	0.020
Max number of homes bid for	J	5
Scale parameter on outside utility	$\natural$	0.9
Interest rate	r	4.0%
LTV requirement	θ^{ltv}	2.0%
Standard deviation of taste shocks	σ_ε	40,000
	Sellers	
Periods observed when setting prices	ς	6
Monthly cost of being on market	c^s	4,000
Price adjustment coefficient	γ^d	0.003
Markup choice set	$\gamma^{\mathbf{u}}$	{-0.03, -0.01, ..., 0.10}
Greediness parameter	β	2.5
Probability of random experimentation	ω	0.03

Notes: Prices and incomes are in DKK-2011.

of the most recent spell of interest-only payments in case any such were made. Combining this register with data on household characteristics yields the subset of mortgage borrowers in the CMA. The age dependent probability of switching to interest-only payments is then given by the share of households at each age who defer installments.

Remaining parameters

The distribution of home quality used for the simulation is the log-normal distribution fitted to the sample of empirical qualities presented in Table 1 of Section 2. The preference parameter on housing is set to 0.3 as the budget share of housing expenses have revolved around this value over the last 40 years, see Dam et al. (2011) and Davis and Ortalo-Magné (2011). The minimal private consumption c^{min} is based on the guidelines for issuing loans by the Danish FSA to mortgage institutions, The Danish Financial Supervisory Agency (2018). The FSA regards a family of two adults and one child to be a low quality borrower if they do not enjoy at least 11,000 DKK worth of monthly consumption (DKK-2018), while a single person may do with 5,000 DKK. Loans not meeting this requirement are flagged as high risk during revision. Based on the composition of single and family households, I set the minimal consumption requirement to 9,000 (in DKK-2011).

There is not much publicly available data on minimum LTV requirements by private banks before the crisis. The Danish mortgage system only allows borrowers to use 80% of the assessed value of a home as collateral for mortgage debt. The remaining debt needed to buy a home is provided as a regular bank loan (or an informal loan by family members) and it was then at the discretion of banks to decide how much down payment to require. However, the FSA enforced a privately sponsored down payment of 5% in 2015 as part of a macro prudential program. The LTV requirement is therefore here set to 2%. Finally, I set the monthly cost of being on the market to 4,000 DKK and the standard deviation of taste shocks to 40,000.

5 Equilibrium dynamics

A key aspect of the model is that there are significant spill-over effects between households acting on the housing market. On the seller side, new entrants on the market learn from the outcomes of previous sellers when setting asking prices. On the buyer side, as will be discussed, placing a bid on a home potentially bars another household from purchasing it. So long as the key parameters shaping the economic environment stay unchanged, we should expect that this market eventually settles into a statistical equilibrium, see Grazzini and Richiardi (2015). By this is meant a ”rest point”, at which the simulation produces stationary time series of its endogenous variables. It implies that the behavior of the population of agents is repeating itself and is as such not changing in a fundamental way from that point on. From even a moderately complex model such as the one at hand, it is possible to construct a swath of variables that evolve endogenously, so whatever set of variables are considered by the modeler when testing for equilibrium, the choice must be guided through intuition and theory. For reasons that will be clear in the following, the variables that is used to test for equilibrium is the stability of the implied positive assortive matching (PAM), the average marginal price of quality, the average selling time and the average fitness of each mark up strategy.

5.1 Equilibrium in the cross section

The model considered in this paper is closely related to an assignment model of the housing market, as investigated by Määttänen and Terviö (2014) and Landvoigt et al. (2015). An assignment model of housing is solved by assuming that n indivisible homes from a continuous quality distribution must be assigned to n households from a continuous income distribution. Homes are traded on a frictionless market where buyers take prices for given. Määttänen and Terviö (2014) show that if preferences are homogeneous, the diminishing marginal rate of substitution between consumption and quality when income increases implies that there will be PAM.

The equilibrium price function of an assignment model ensures that the first order conditions of all households are fulfilled. Given PAM, this means that the marginal price of quality at any point in the quality distribution depends on the marginal rate of substitution between quality and consumption of the household at the same quantile in the income distribution. This, in turn, means that prices depend on the relative spread in quality and income distributions. If, at any quantile k, there is a high dispersion in income and a lower dispersion in quality, the result is a steeper increase in the price function. This is because the marginal buyer just above k will not experience a change in available quality corresponding to the change in income from k, which entails higher marginal rate of substitution and therefore higher prices in equilibrium. A notable conclusion from this reasoning is that if the quality distribution is a scaled version of the income distribution, then prices increase linearly with quality, see Landvoigt et al. (2015).

In the context of the model at hand, this would translate into the condition that along the price function, there is a buyer with an income y and wealth w who chooses to bid for a home of quality q because of the optimality condition

$$p'(q) = \frac{y - \xi(r)(p(q) - w)}{q} \frac{\alpha}{1 - \alpha} \tag{30}$$

is fulfilled. Solving the assignment model mentioned above comes down to finding a price function such that the optimality condition above obtains for all households in the period considered. In both papers, it requires a boundary condition that the worst home has price p_0. The first order condition of buyers then yield a differential equation which is solved to get prices. The upshot is that changes in the willingness to pay in the bottom of the income trickles all the way up through the price function affecting all prices above.

These points also underline the way our present simulation departs from an assignment model. Because of the absence of explicit market clearing conditions, equation (30) is

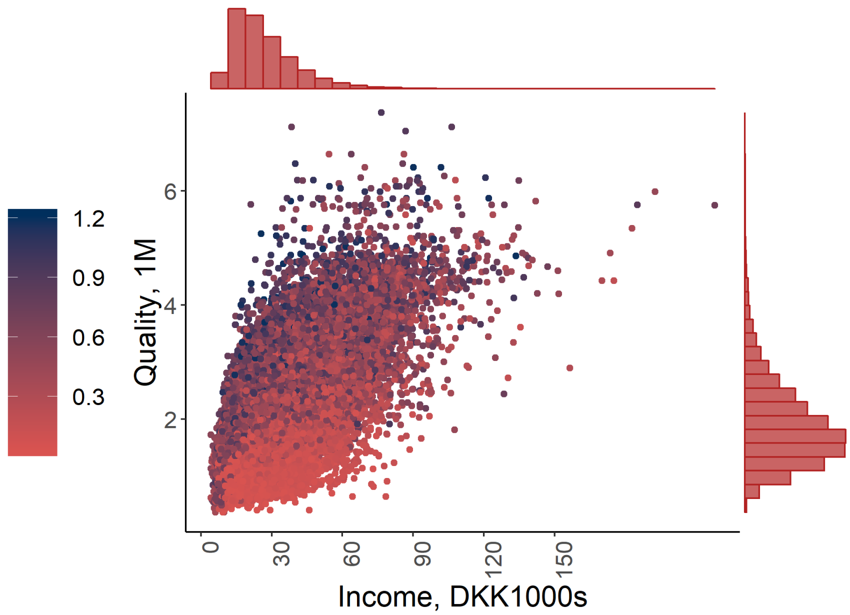

Figure 4: Positive Assortive Matching among households.

Notes: Each dot represents a household. Dots are colored according to the home equity of households, measured in millions DKK, see left legend. The density of monthly household income is displayed at the top of the figure frame and the density of home quality is at the right of the frame.

only (approximately) fulfilled along the price function when the model finds a stationary equilibrium. It is in particular not fulfilled the moment willingness to pay among buyers shifts because of, e.g., an interest hike. The marginal *asking* price for quality stays unchanged in the period of a shock due to the backward looking expectation formation of prices, the learning mechanism of markup setting and the downward rigidity in markups.

An interest hike will therefore cause buyers to bid for homes in a lower segment of the quality distribution than before, creating excess demand and thus changing the liquidity of homes along the quality distribution. Through this change of liquidity, prices adjust downward in the upper segments and upward in the lower, since a higher density of buyers for a segment increases the probability that a seller meets a buyer willing to accept a positive markup.

Getting at such effects is the very point of applying a simulation approach to a situation like the recent bust cycle where housing liquidity changed dramatically.

5.2 Stationarity of outcomes

Table 3 shows the results of Augmented Dickey-Fuller (ADF) tests for stationarity on the above mentioned variables. The first block yields a test for the stationarity of the pattern of PAM in the simulation. In each period of the simulation after a suitable burn in phase, the rank of the quality of household i's home is regressed on its rank in the income and wealth distribution. Storing the coefficients from each of these regressions produces a set a time series on which the ADF test is applied. Evidently, they show that the degree of sorting into quality is stationary throughout the simulation as the null of a unit root is rejected.

The second block of Table 3 test variables that are informative of market conditions. As noted in the model section, a grid of centroids, that each contain a LOESS regression of price on quality, is distributed evenly in the support of the quality distribution. In order to aggregate the information contained in the centroids, the average over intercepts and coefficients are taken at each time t. This again forms a time series which shows no unit root. The same goes for the period-wise average time it takes to sell homes. The last block shows the ADF test performed on the series of period-wise average fitness of each initial markup, which also clearly displays stationarity.

To get at feeling for the behavior of these time series, Figure 13 in Appendix C displays a range of average statistics from simulation that produced time series for the stationarity tests.

5.3 Concave demand and price momentum

Within the housing search literature, an extensive effort has gone into explaining why housing prices feature a large degree of momentum. A compelling explanation that may be embedded in a search framework which was made recently by Guren (2018), is that sellers are facing a concave demand curve for their homes.[10] He finds empirically, that the probability of

[10] This is a recent paper, but it relies on the notion of kinked demand curves being responsible for price rigidities, which has been around since the 1930s; Dupraz (2017)

Table 3: Augmented Dickey-Fuller test for stationarity. Alternative hypothesis is stationarity.

Variable	**P-value of ADF**	**Mean**	**Std. err**
	Positive assortive matching		
$Rank(q_{it}) = \beta_{inc} Rank(y_{it}) + \beta_w Rank(w_{it}) + \varepsilon_{it}$			
β_{inc}	0.049	0.6460	0.003
β_w	0.045	0.3610	0.003
$\beta_{inc} + \beta_w$	$\leq$ 0.019	1.0070	
	Market conditions		
Intercept, prices	$\leq$ 0.010	271,117	(12209.855)
Mean gradient, prices	$\leq$ 0.143	0.739	(0.010)
Time-on-market	$\leq$ 0.010	2.771	(0.183)
	Markup fitness		
Initial markup			
-0.01	$\leq$ 0.01		
0.00	$\leq$ 0.01		
0.01	$\leq$ 0.01		
0.02	$\leq$ 0.01		
0.03	$\leq$ 0.01		
0.04	$\leq$ 0.01		
0.05	$\leq$ 0.01		
0.06	$\leq$ 0.01		
0.07	$\leq$ 0.01		
0.08	$\leq$ 0.01		
0.09	$\leq$ 0.01		

selling a home is concave in it's relative price to other similar homes. This implies a strategic complementarity in pricing setting where raising one's price above average quickly reduces the chances of selling, yet undercutting the market is not similarly effective in alluring buyers. The reason for this phenomenon is, presumably, that buying a home comes in a two step procedure; first buyers must inspect a home (which is costly), and then upon inspection decide whether to purchase the home. Based on the data in Guren (2018), it appears that setting one's home at market price is sufficient to get enough buyers through the door to ensure a high probability of selling, so there is no benefit in cutting prices below market level.

It follows from the strategic complementarity of this setting that a rational seller who realizes that fundamentals have just changed will be reluctant to set a new price that deviates much from that of others. It is however shown theoretically by Guren (2018) that concave demand in itself does not in itself create price rigidity, it merely works to amplify it significantly.[11] He therefore adds extra rigidity in the form of either staggered price setting or backward looking sellers. A certain share of of non-rational sellers therefore works as a slow-moving anchor that forces rational agents to only update prices sluggishly to new fundamentals.

Note here that we should also expect the notion of concave demand and strategic complementarity to apply in the current context. The reason is that buyers have extensive knowledge of the menu of homes for sale, but were only able to submit J bids pr period. Since the number of buyers is finite, a zero sum game of garnering bidders emerges between sellers because the more interested buyers a seller faces, the better a chance that she will get to sell to one of them. (Recall that after bids were submitted, sellers were selected in random order to accept the highest they were offered. Sellers with few bids, as well as the last in line, therefore faced the chance of loosing all their bidders.) As in Guren's model, there is a significant difference in the probability of selling when going from seeing 1 bid for one's home to 2 or 3. However, once a sufficient number of bids have been obtained for a home due to

[11] Intuitively, if all sellers are rational then everyone will update prices symmetrically upon a realized shock and no momentum emerges.

favorable pricing, chances are that at least one of the buyers will end up purchasing it and so the marginal effect of lowering the price further is minimal.

Figure 5: Concave demand at low markups.

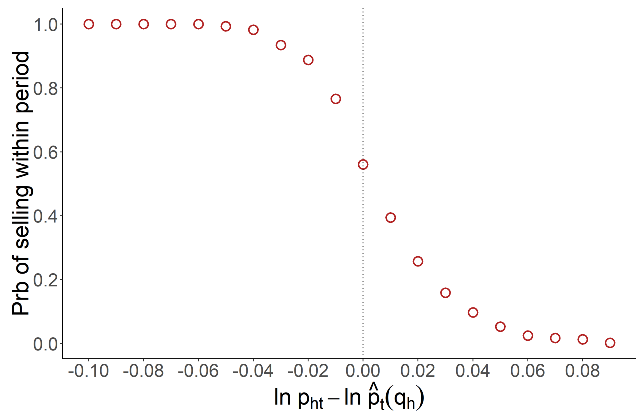

Figure 5 shows that this pattern indeed arises in the simulation. It displays the fraction of homes that are sold within a period given their markup in bins of 1 percent, averaged over 1000 periods in which no shocks occur. Note that the current markup of home h, as defined in equation (11) and (12), is given by the difference between its asking price and the current estimate of the market price of h's quality, $\ln p_{ht} - \ln \hat{p}_t(q_h)$. We see in the figure that there is a gain in probability of selling when cutting prices 1% below the expected price, but the marginal effect diminishes after that. Of course, to a certain extent the diminishing marginal effects comes from closing in on the upper bound of 1, but note that the selling probability could theoretically have jumped to 1 after the 1% negative markup, if decreasing marginal effects were not the case. A further note is that the shape of the curve is partly affected by the variance of idiosyncratic taste shocks as a larger variance increases the spread in willingness to pay for a particular home.

As in Guren (2018), the simulation features a significant degree of price momentum. Different from the model at hand, that setup features a set of forward-looking rational agents whose price setting is kept in check by backward looking price setters. In my model, all agents are by default backward-looking, yet some are selected in each period to experiment with markups so as to reveal if any strategies have changed in profitability due to a change in conditions. The fact that changing conditions have to be learned through experimentation generates rigidity in prices on its own. Consider for example the case where the interest rate falls, increasing the purchasing power of home buyers. The resulting income effect works to decrease buyers' sensitivity to markups, which result in increased sales of homes that carry a high markup. These sales push up asking prices in the periods following as new sellers entering the market observe higher average selling prices together with increased fitness of high markups. Followingly, they have higher expectations of sales prices and fitness of high markups, causing them to increase asking prices further. This process of adaptive groping stops to increase prices when buyers are no longer willing to accept positive markups in the average. This process is time consuming and therefore leading to momentum in prices.

As was noted above, the momentum from learning is further amplified by the strategic complementarity that the market clearing mechanism induces. A single seller who tries to set a price at the new higher steady state level, immediately after fundamentals have changed, will not succeed in selling if this is too far away from the present level as buyers will substitute away from this particular home. The learning mechanism, in conjunction with costly bidding, is therefore what drives momentum in the simulation. We will see the presence of momentum in the experiments with the model.

6 Housing bust

We now turn to the primary experiment with the model, creating a bust in housing prices resembling that of the late 2000s in Copenhagen. Given that the important parameters of the model are not directly fitted to data as of now, its predictions are thought of as qualitative, and more work on the calibration is needed to replicate the data in a precise way. Two factors are used to generate a swing of the same character as we saw in the introduction; (i) the effective interest rate on new mortgages tops in July 2008 and then drops sharply. (ii) The spiking interest rate is accompanied by falling prices outside Copenhagen. The first effect (i)

is implemented as continuous increase in r_t on new mortgages by 9.4 basis points monthly from 4% in t_0 to 7% in t_{32}. After t_{32} the mortgage keeps falling until it reaches 3% in t_{56}, after which it decreases by 1 pct point over the following 60 months. This trajectory roughly imitates the actual interest rate development from 2006 to 2016 depicted in Figure 2. Effect (ii) is given by a 30% gradual fall in the marginal price of quality in the outside option, p^o. Like the interest hike, this fall is evenly distributed over the period t_0-t_{32} and then starts rebounding gradually in the period t_{56}-t_{116}.

The effect of changing the interest rate was touched upon in Section 5.1. As financing costs go up, the marginal rate of substitution between housing and consumption goes down, essentially creating a flatter price function. The flattening occurs when homes in the top quality segments go unsold, causing sellers to downgrade asking prices until the point when they are able to attract buying households of the top income/wealth segments again. Meanwhile, cheap homes in the bottom quality segment are in relatively higher demand, causing them to sell faster and at higher markups. A visualization of this effect can be found in Figure 10 of Appendix C. Here, the interest rate follows the trajectory described above, while prices in the outside option remain fixed at previous levels. The result of the interest rate hike is that low quality homes actually increase slightly in price while higher quality homes plummet. The pattern is reversed as interest rates start falling again.

This behavior of prices along segments is somewhat extreme because it partly stems from forcing households to buy and sell after receiving a moving shocks, λ^f, even if that was welfare improving for them to do so. Those households are typically going for homes in the bottom segment, which keeps prices there afloat. We should therefore minimize the prevalence of forced movers during the crisis while avoiding that the market goes into a halt. I thus set $\lambda^f = 0.05$ during the bust period to ensure that the majority of households going on the market after the onset of the crash did so expecting to make a utility gain.

When, in addition, outside prices are set to fall together with prices inside the city, things look very different. As outside option prices fall, their value as a threat point increases,

Figure 6: Reaction in housing market liquidity and price by segment from an initial interest rate hike and a following interest rate drop.

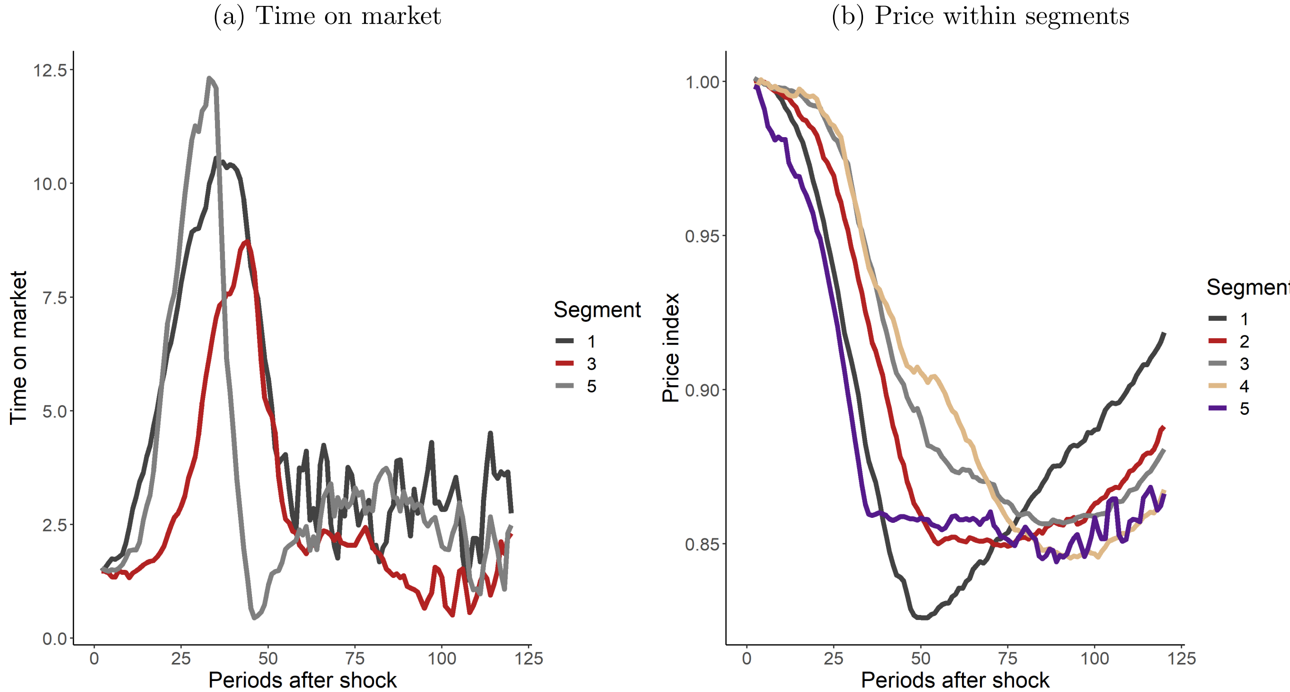

causing buyers to submit lower bids than before. The result is that more homes than before do not receive any bids above their posted price, implying they will go unsold until next period. The drop in outside prices puts downward pressure on homes of all quality levels since homes of all qualities are available in the outside option.

We notice in Panel (b) of Figure 6 that the impact is not uniform across quality segments. The bottom segment immediately takes the most drastic hit together with the top segment, while the middle segments fall less. The drop in prices is accompanied by corresponding increases in time-on-market along segments, shown in Panel (a) Figure 6. The excessive price fall in the top housing segment compared to the middle segments is to be expected when all households in the top income segments downgrade their demand for quality, leaving no demand for high quality homes at the current asking prices. Households in the middle income segments also downgrade their demand for quality, but their demand for for mid-quality homes is more than replaced by demand from top earners. Hence, mid-quality homes fall less than top quality.

Low quality home prices fall dramatically in the beginning of the bust because low income households may opt out of the market entirely. The dynamics of the bust is such that falling prices erode save up home equity. Falling home equity implies accelerating financing costs since homes can only be collateralized for mortgages up to 80%. Higher leveraging must be financed with an expensive bank loan and this hurts in particular buyers of low quality homes, who are often both low on income and saved up equity. Increasing interest rates and falling prices therefore cause a strong contraction in the low quality segment.

The interest rate hike ends in period 32 at which point the lowest quality segment immediately stops falling further. Low quality homes get in relative high demand as soon as borrowing conditions stabilize since voluntary buyers no longer opt out of the market, although facing increased financing costs and eroded home equity. Thus they direct their demand to this segment. Prices in the other segments keep falling several periods after the interest hike has ended because they are still not sufficiently low to attract buyers due to the inherent momentum of the adaptive price setting mechanism discussed in the previous section.

Around period 50, prices in the top segment start to pick up as the bulk of available homes has been sold off. This is expressed by the fact that average time-on-market is now back to an average of around 3 months, see Panel (a) of Figure 6. In the periods following, prices start increasing at a roughly equal pace along segments, although at different points in time. The top segment is the first to start rebounding, then comes the second segment, the third, etc. The result is that the five price indices are ordered according to their quality segmentation during the rebound.

As it stands, the model therefore misses the empirical fact, that prices in the top segment fell the least but does get price behavior in the recovery phase right. Future work on the model will seek to deal with this. One reason why the model currently cannot replicate the empirical price movements in the bust is that the mortgage rate offered is symmetrical for low and high income households. Relaxing this condition will change the relative demand for high and low quality homes in favor of the top segment, bringing model behavior closer to data. Another reason may be that resident households are the only actors on the market in the model. Wealthy outside individuals and institutional investors can be important for the development in prices and they tend to favor high end homes, see Cvijanovic and Spaenjers (2020) and Favilukis and Van Nieuwerburgh (2017). Including non-resident investors in the model may be important for the behavior of prices during a bust.

Coming to the comparison with the empirical magnitudes reviewed in Section 2, consider first the simulated drop in prices. This amounts to about 15% which, although sizeable, is still about 15 percentage points short of the drop observed in data between 2006 and 2009. This is not entirely surprising as the model has not taken the increased unemployment rate and the change in consumer expectations into account. Given that the reaction in prices is less than empirically observed, the associated time-on-market is a bit on the high-side. Sales time in the simulation increases during the bust from 2 months to 8-12.5 months, depending on segment, while the empirical counterpart for the CMA was rather a rise from 3 to 7 months.

7 Conclusion

The Copenhagen Metro Area experienced a dramatic boom-bust(-boom) cycle in the years since 2003. Interestingly, while home prices escalated uniformly across housing segments in the years of the boom, the response was much less symmetrical as the bust set in.

This paper has developed a simulation approach to understand why such a pattern occurred. Several factors have likely contributed to the event, but the analysis here focuses on the impact of the effective mortgage rate. The simulated model underlines the importance of heterogeneity among both households and homes for an understanding of housing price fluctuations. In short, when financing costs changes, marginal prices along the distribution of home quality changes because of local changes in demand.

The model also creates a tight link between prices and market liquidity in that (*i*) posting a price above buyers' reservation price results in longer time to sell and (*ii*) a contraction in the volume of buyers reduces the probability of selling putting downward pressure on prices. Temporary overpricing occurs, for example when credit conditions changes, due to a continuous learning process among sellers. Large negative corrections in posted prices were indeed observed during the housing bust, serving the assumption that sellers were continuously updating their view on the market value of their homes. In effect, the simulation therefore shows a sharp decrease in market liquidity during the bust as observed.

Policy makers often care as much about cross-sectional differences as average effects and therein lies the strength of the simulation approach taken in this paper. It proves very informative about the dynamics of cross-sectional effects in face of a shock to market conditions or policy. Whereas the simulation returns predictions qualitatively on par with observed data, a finer calibration or estimation is needed to make predictions for policy purposes. This will be a venture for future work. Yet given the current level of calibration, results are encouragingly reasonable.

References

Anufriev, M., Bao, T., and Tuinstra, J. (2016). Microfoundations for Switching Behavior in Heterogeneous Agent models: An Experiment. *Journal of Economic Behavior & Organization*, 129:74–99.

Arifovic, J. and Karaivanov, A. (2010). Learning by doing vs. learning from others in a principal-agent model. *Journal of Economic dynamics and control*, 34(10):1967–1992.

Bao, T. and Hommes, C. (2015). When speculators meet constructors: positive versus negative feedback in experimental housing markets. *CeNDEF working paper*.

Baptista, R., Farmer, J. D., Hinterschweiger, M., Low, K., Tang, D., and Uluc, A. (2016). Macroprudential policy in an agent-based model of the uk housing market. Technical Report 619, Bank of England.

Berger, D., Guerrieri, V., Lorenzoni, G., and Vavra, J. (2017). House prices and consumer spending. *The Review of Economic Studies*, 85(3):1502–1542.

Case, K. E., Shiller, R. J., and Thompson, A. (2012). What Have They Been Thinking? Home Buyer Behavior in Hot and Cold Markets. Technical report, National Bureau of Economic Research.

Cvijanovic, D. and Spaenjers, C. (2020). 'We'll Always Have Paris': Out-of-Country Buyers in the Housing Market. SSRN Scholarly Paper ID 3248902, Social Science Research Network, Rochester, NY.

Dam, N. A., Hvolbøl, T. S., Pedersen, E. H., Sørensen, P. B., and Thamsborg, S. H. (2011). Udviklingen på ejerboligmarkedet i de seneste år - Kan boligpriserne forklares? *Technical report, Central Bank of Denmark*.

Davis, M. A. and Ortalo-Magné, F. (2011). Household expenditures, wages, rents. *Review of Economic Dynamics*, 14(2):248–261.

Davis, M. A. and Van Nieuwerburgh, S. (2015). Chapter 12 - Housing, Finance, and the Macroeconomy. In Duranton, G., Henderson, J. V., and Strange, W. C., editors, *Handbook of Regional and Urban Economics*, volume 5 of *Handbook of Regional and Urban Economics*, pages 753–811. Elsevier.

Dupraz, S. (2017). A Kinked-Demand Theory of Price Rigidity. *Mimeo*.

Favilukis, J., Ludvigson, S. C., and Van Nieuwerburgh, S. (2017). The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium. *Journal of Political Economy*, 125(1):140–223.

Favilukis, J. and Van Nieuwerburgh, S. (2017). Out-of-Town Home Buyers and City Welfare. *SSRN Electronic Journal*.

Geanakoplos, J., Axtell, R., Farmer, J. D., Howitt, P., Conlee, B., Goldstein, J., Hendrey, M., Palmer, N. M., and Yang, C.-Y. (2012). Getting at Systemic Risk via an Agent-based Model of the Housing Market. *American Economic Review*, 102(3):53–58.

Gintis, H. (2007). The Dynamics of General Equilibrium. *The Economic Journal*, 117(523):1280–1309.

Glaeser, E. L. and Nathanson, C. G. (2015). An Extrapolative Model of House Price Dynamics. Working Paper 21037, National Bureau of Economic Research.

Grazzini, J. and Richiardi, M. (2015). Estimation of Ergodic Agent-based Models by Simulated Minimum Distance. *Journal of Economic Dynamics and Control*, 51:148–165.

Guren, A. M. (2018). House Price Momentum and Strategic Complementarity. *Journal of Political Economy*, 126(3):1172–1218.

Haldrup, K., Staunstrup, J. K., Hansen, J. Z., Jørgensen, E. H., and Stephensen, P. (2015). Tvangsauktioner i Danmark. Technical report, BVC, Aalborg University, DREAM.

Han, L. and Strange, W. C. (2015). The Microstructure of Housing Markets: Search, Bargaining, and Brokerage. In *Handbook of regional and urban economics*, volume 5, pages 813–886. Elsevier.

Heathcote, J., Perri, F., and Violante, G. L. (2010). Unequal we stand: An empirical analysis of economic inequality in the united states, 1967–2006. *Review of Economic dynamics*, 13(1):15–51.

Hommes, C. (2011). The heterogeneous expectations hypothesis: Some evidence from the lab. *Journal of Economic dynamics and control*, 35(1):1–24.

Khandani, A. E., Lo, A. W., and Merton, R. C. (2013). Systemic risk and the refinancing ratchet effect. *Journal of Financial Economics*, 108(1):29–45.

Kuminoff, N. V. and Jarrah, A. S. (2010). A new approach to computing hedonic equilibria and investigating the properties of locational sorting models. *Journal of Urban Economics*, 67(3):322–335.

Kuminoff, N. V., Smith, V. K., and Timmins, C. (2013). The new economics of equilibrium sorting and policy evaluation using housing markets. *Journal of Economic Literature*, 51(4):1007–62.

Landvoigt, T., Piazzesi, M., and Schneider, M. (2015). The Housing Market (s) of San Diego. *American Economic Review*, 105(4):1371–1407.

Merlo, A., Ortalo-Magné, F., and Rust, J. (2015). The Home Selling Problem: Theory and Evidence. *International Economic Review*, 56(2):457–484.

Moen, E. R., Nenov, P., and Sniekers, F. (2015). Buying first or selling first in housing markets. *Mimeo*.

Määttänen, N. and Terviö, M. (2014). Income Distribution and Housing Prices: An Assignment Model Approach. *Journal of Economic Theory*, 151:381–410.

Mulalic, I., Rasmussen, H., Rouwendal, J., and Woltmann, H. H. (2017). The Financial Crisis and Diverging House Prices: Evidence from the Copenhagen Metropolitan Area. *SSRN Electronic Journal*.

Piazzesi, M. and Schneider, M. (2016). Housing and Macroeconomics. In Taylor, J. B. and Uhlig, H., editors, *Handbook of Macroeconomics*, volume 2, pages 1547–1640. Elsevier.

Rosen, S. (1974). Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition. *Journal of Political Economy*, 82(1):34–55.

Salle, I., Yildizoglu, M., Zumpe, M., and Sénégas, M.-A. (2017). Coordination through social learning in a general equilibrium model. *Journal of Economic Behavior & Organization*, 141:64–82.

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological review*, 63(2):129.

Sinitskaya, E. and Tesfatsion, L. (2015). Macroeconomies as constructively rational games. *Journal of Economic Dynamics and Control*, 61:152–182.

The Danish Financial Supervisory Agency (2018). Vejledning til regnskabsindberetning for kreditinstitutter og fondsmæglerselskaber m.fl.

Appendix A

Micro data on housing transactions

From Statistics Denmark, the sales register EJSA is obtained and merged with a detailed register of homes, BOL, to create a data set of all sold freeholds in the Copenhagen Metro Area between 2003-2015.[12] Each sales record in EJSA carries information about transaction price, the date of transaction, a unique address identifier and the type of the transaction.

The BOL register contains a host of features for each home, including indoor size in square meters, housing type, year of construction, and an address identifier that matches the identifier of EJSA. In order to determine a geographic location for each home, Statistics Denmark has provided the commuting zone of each address. Commuting zones are based on the National Traffic Model that measures commuting times and congestion. They are relatively fine grained in densely inhabited areas with a median size of 2 km^2 in the CMA. Homes in the final data set for the CMA are distributed across 129 different commute zones.

Combining EJSA, BOL and commuting zones yields a data set of 137,700 sales between 2003-2015 when conditioning on the following: square meters of indoor space is in the range [30, 300], deflated transaction prices (in DKK-2015) are in the range [0.3, 12] million DKK, that commute zones, transaction date, housing type and construction year are all available, that the traded home must be a freehold designated for all-year living, that there must be at least 10 observed trades pr included commute zone, and the trade must carried out on free market conditions (i.e. not traded between family members).

The remaining data applied in Section 2 is publicly available.

[12]The set of municipalities that define Copenhagen Metro Area in this paper is Copenhagen, Frederiksberg, Ballerup, Brøndby, Dragør, Gentofte, Gladsaxe, Glostrup, Herlev, Albertslund, Hvidovre, Høje-Taastrup, Rødovre, Ishøj, Tårnby, Vallensbæk.

Appendix B

Housing prices and regression parameters

Figure 7: Price of quality in the Greater Copenhagen Area in selected years.

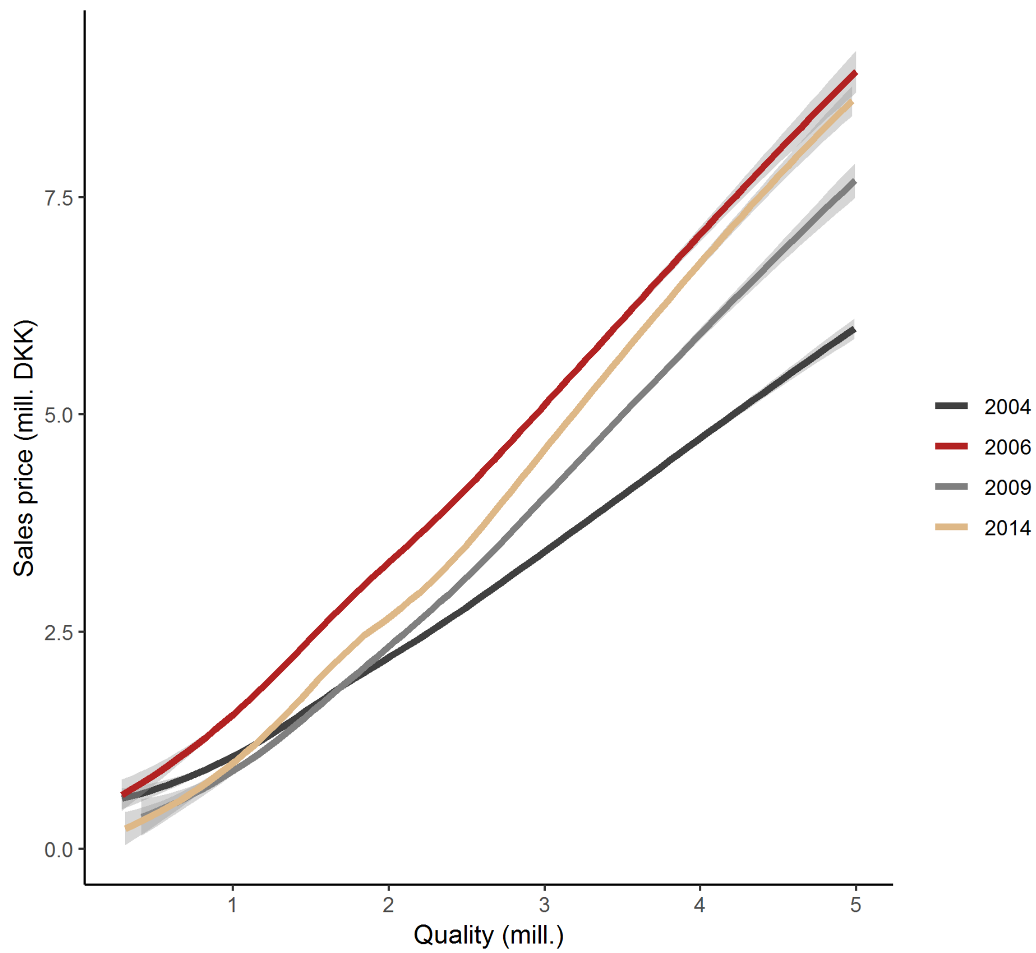

Notes: A LOESS kernel for the relation between imputed quality and real sales price has been estimated for each year. Quality is imputed as in equation (2).

Figure 8: Rank of homes in quality distribution in 2003 and 2010

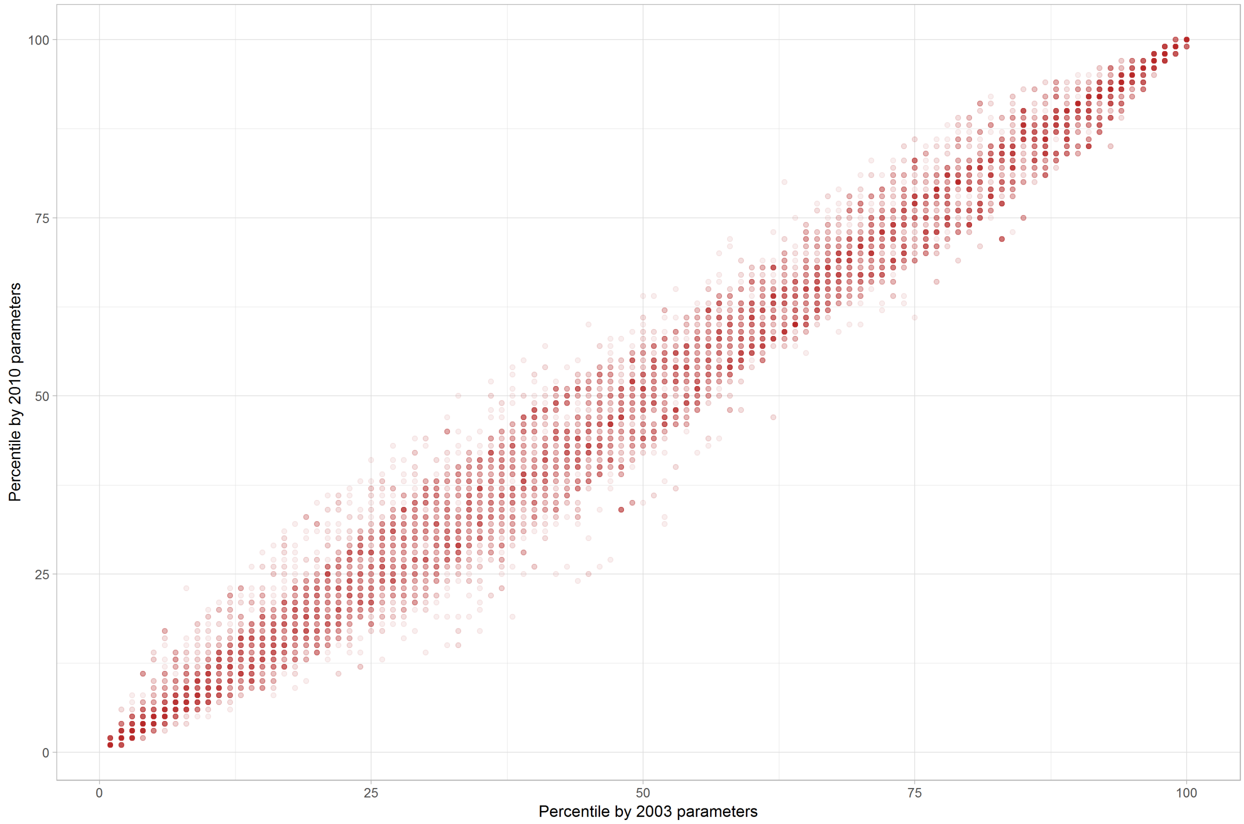

Notes: Each dot represents a home sold in 2003. Its x-coordinate represents its percentile in the quality distribution by the hedonic regression described in section 7. Y-coordinates represent the percentile in the quality distribution of the same home if using coefficients from the same hedonic regression applied to sales in 2010 instead of 2003.

Table 4: Estimates of hedonic regression in equation (1)

Regressor	Estimate	Std. Error	Regressor	Estimate	Std. Error
log.size	0.416598	0.004300	ltm147220	12.679526	0.029372
dhouse	0.232714	0.007498	ltm147230	12.659355	0.027833
d1900	0.110741	0.009874	ltm147240	12.815717	0.055177
d00_50s	0.042963	0.006049	ltm147250	12.709605	0.038892
d90_00s	0.075236	0.016518	ltm151010	12.384216	0.028617
ltm102120	12.854134	0.039234	ltm151040	12.464675	0.058423

ltm102140	12.778117	0.041363	ltm151060	12.389633	0.034742
ltm102150	12.750958	0.042920	ltm151070	12.441693	0.030331
ltm102160	12.612236	0.061922	ltm151080	12.424124	0.031306
ltm102170	12.723933	0.075896	ltm153010	12.502430	0.044436
ltm102180	12.551194	0.045196	ltm153020	12.383304	0.028520
ltm102210	12.817983	0.030596	ltm153030	12.219505	0.030799
ltm102220	12.865168	0.032283	ltm155010	12.552439	0.025130
ltm102230	12.703245	0.049421	ltm155020	12.495549	0.054710
ltm102310	12.730348	0.035614	ltm157110	13.142776	0.051271
ltm102320	12.679431	0.025364	ltm157120	12.818769	0.022639
ltm102330	13.015512	0.048070	ltm157130	12.861886	0.029408
ltm102340	12.671220	0.021531	ltm157140	12.954397	0.027035
ltm102350	12.678811	0.025733	ltm157150	12.717015	0.032871
ltm102410	12.587627	0.024500	ltm157210	12.956260	0.029230
ltm102420	12.830067	0.075587	ltm157220	12.697402	0.036039
ltm102430	12.540635	0.027023	ltm157230	12.825981	0.032418
ltm102440	12.634756	0.023713	ltm157240	12.642455	0.025810
ltm102450	12.534150	0.021095	ltm157250	12.563862	0.030541
ltm102510	12.673716	0.059686	ltm157260	12.548395	0.030082
ltm102520	12.570961	0.039943	ltm159010	12.550124	0.032581
ltm102530	12.515309	0.026810	ltm159020	12.391372	0.033911
ltm102540	12.490196	0.023309	ltm159030	12.572960	0.030379
ltm102550	12.513860	0.029293	ltm159040	12.521818	0.037589
ltm102560	12.519132	0.030508	ltm159050	12.389545	0.051056
ltm102610	12.532847	0.024278	ltm159060	12.548221	0.029059
ltm102620	12.500427	0.027191	ltm159070	12.559003	0.023194
ltm102630	12.520444	0.022456	ltm159080	12.602388	0.023999

ltm102640	12.557618	0.026848	ltm159090	12.610865	0.040625
ltm102650	12.465345	0.038999	ltm161010	12.397824	0.045802
ltm102660	12.512393	0.029194	ltm161020	12.363884	0.030424
ltm102710	12.532439	0.038858	ltm161030	12.372594	0.028453
ltm102720	12.549216	0.023949	ltm163010	12.443592	0.035880
ltm102730	12.515972	0.034870	ltm163020	12.455548	0.031132
ltm102740	12.668465	0.057025	ltm163030	12.452047	0.037300
ltm102750	12.448391	0.029664	ltm165010	12.331571	0.025661
ltm102760	12.500660	0.031851	ltm165040	12.464866	0.040683
ltm102770	12.525895	0.026732	ltm167010	12.432137	0.022835
ltm102810	12.731806	0.028125	ltm167020	12.400341	0.031368
ltm102820	12.605513	0.026909	ltm167030	12.422284	0.035178
ltm102840	12.463651	0.045085	ltm167040	12.442020	0.026519
ltm102860	12.492339	0.024927	ltm167050	12.374110	0.036810
ltm103130	13.160196	0.053270	ltm169010	12.396963	0.025608
ltm103140	12.804217	0.026596	ltm169020	12.340871	0.042431
ltm103150	12.577607	0.071743	ltm169030	12.302471	0.048240
ltm103160	12.497259	0.024545	ltm169040	12.344392	0.025332
ltm103170	12.583409	0.021122	ltm169050	12.264261	0.027962
ltm103180	13.003035	0.058375	ltm175010	12.445907	0.034060
ltm103190	12.671534	0.040895	ltm175020	12.488916	0.024204
ltm103210	12.564292	0.025421	ltm175030	12.480945	0.029904
ltm103220	12.510857	0.022380	ltm175040	12.453580	0.033470
ltm103230	12.537008	0.025626	ltm175050	12.393356	0.038264
ltm103240	12.520137	0.025322	ltm183010	12.358173	0.023015
ltm103250	12.475482	0.034006	ltm183020	12.417117	0.058329
ltm103280	12.419685	0.055361	ltm185120	12.545980	0.031530

ltm147110	12.710155	0.025506	ltm185130	12.503484	0.037358
ltm147120	12.729381	0.030451	ltm185140	12.440331	0.031108
ltm147130	12.791478	0.026855	ltm185150	12.415797	0.031819
ltm147140	12.670477	0.046839	ltm185160	12.379719	0.034768
ltm147150	12.719904	0.039888	ltm185170	12.391499	0.027007
ltm147160	12.651240	0.022125	ltm187010	12.338446	0.024770
ltm147210	12.698657	0.028882	ltm187020	12.400770	0.037186

Figure 9: Life-cycle profiles in simulation.

(a) Household income in simulation

(b) Probability of interest-only payments

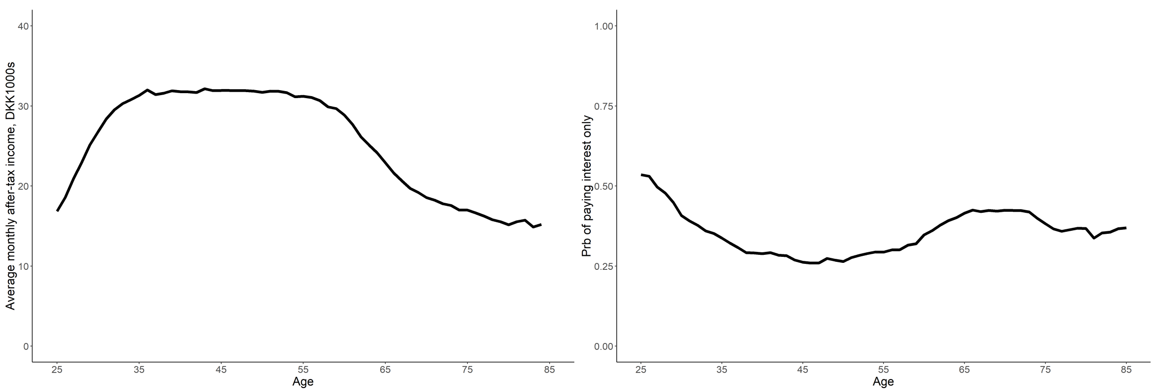

Appendix C

Additional model output

Figure 10: Prices in segments when only the interest rate changes and prices in the outside option are kept fixed.

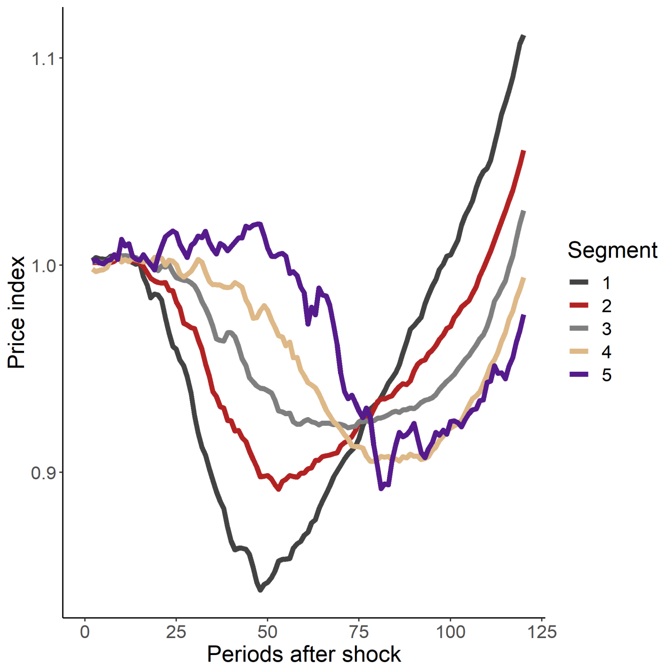

Figure 11: Liquidity effects of bust starting in $t = 800$.

(a) Average loan-to-value

(b) Attentive movers

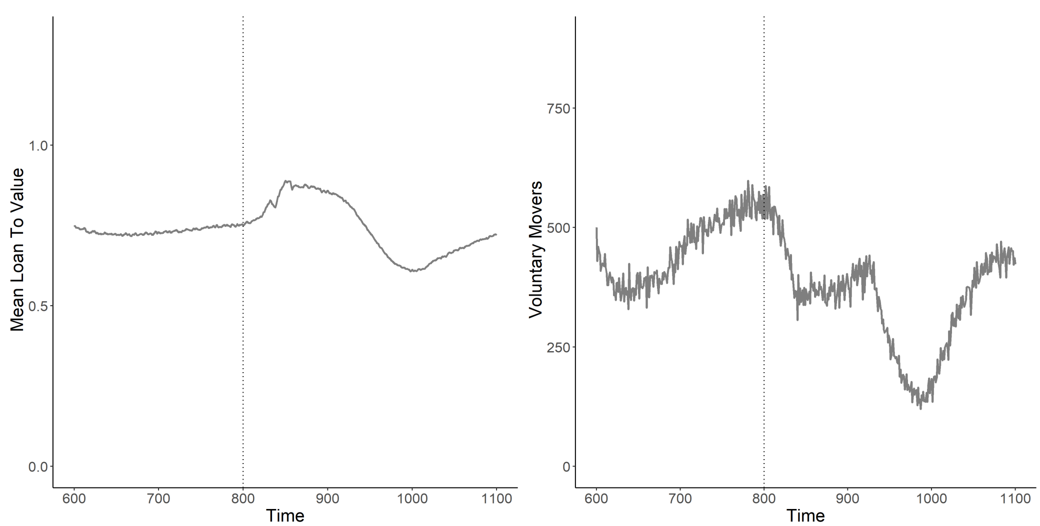

Figure 12: Liquidity effects of bust starting in $t = 800$.

(a) Households constrained from reducing prices

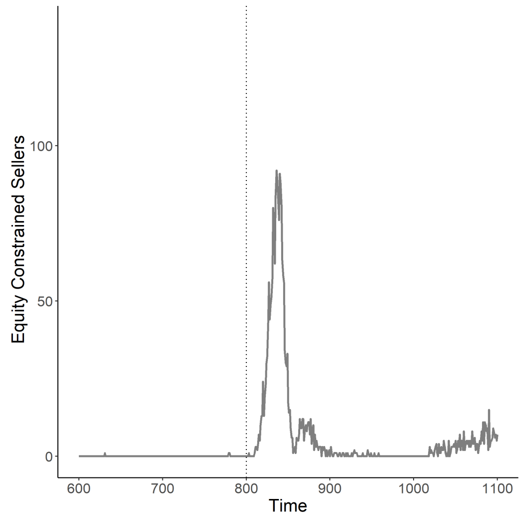

(b) Fitness of markup strategies

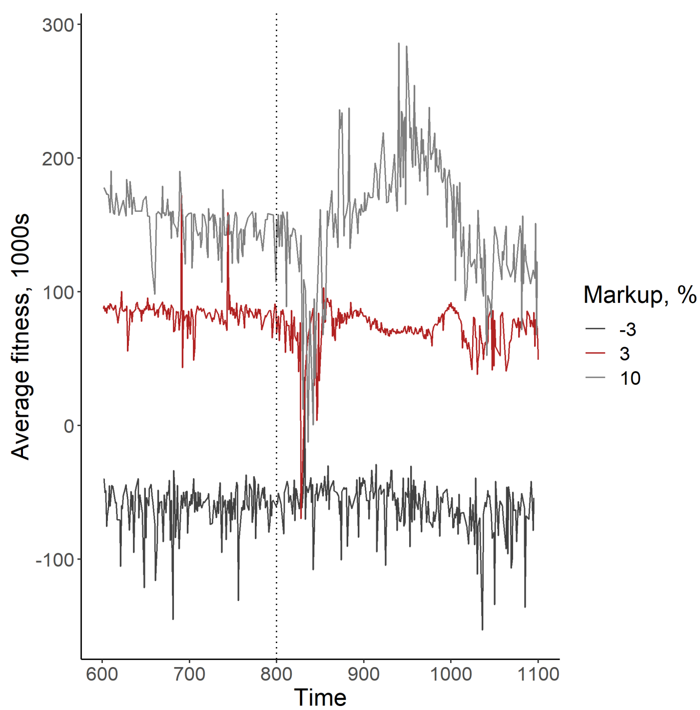

Figure 13: Time series from simulation without shocks.

(a) Time on market

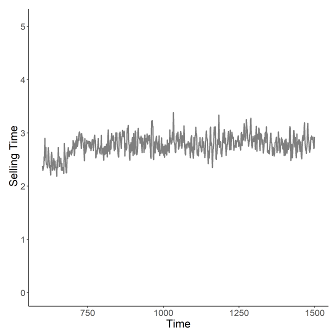

(b) Average marginal price of quality

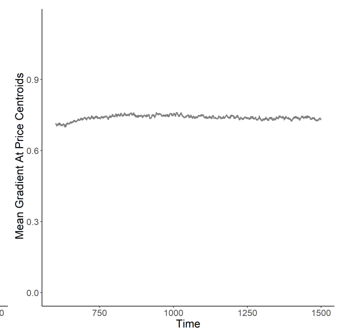

(c) Households not selling due to equity constraint.

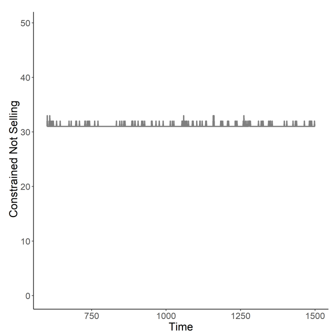

(d) Average LTV

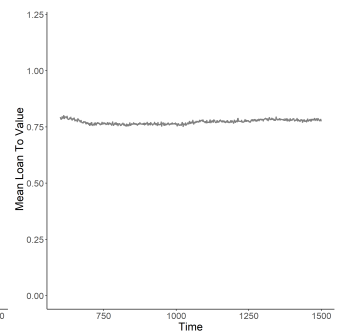

(e) Voluntary movers

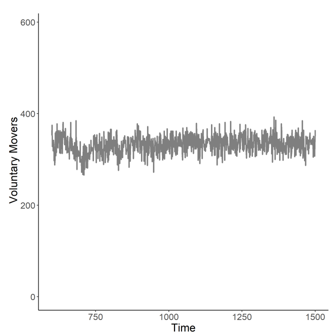

(f) Average fitness of markups

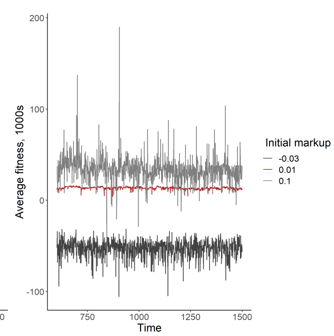

(g) Households paying interest-only on mortgage in t

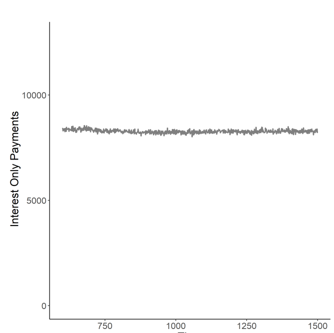

(h) Utility of outside option

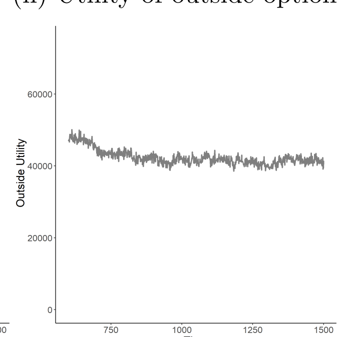

Chapter 2

A Dynamic Equilibrium Model of Commuting, Residential and Work Location Choices

With: Maria Juul-Hansen[‡], Fedor Iskhakov[♯], John Rust[△]
and Bertel Schjerning[†]

In this paper we develop and estimate a dynamic equilibrium life cycle model of residential and work location choices. In our model, commuting is endogenously determined by the distance between work and residence, and house prices are determined in equilibrium. We estimate the model using Danish register data for the entire population of households in the Greater Copenhagen area (GCA). Assuming a fixed supply of housing in the short run, we consider the effects on house prices, job mobility, residential sorting and commuting in two counterfactual equilibria with i) increased supply of housing in the center of the GCA and ii) increased cost of commuting between all residential and work location regions. We find that i) results in lower prices in equilibrium for all regions and a higher degree of urbanization. ii) implies lower average commute times, but also a higher share of people in non-employment, in particular for residents outside of the GCA. The equilibrium prices drop in peripheral regions with low job density.

[‡]University of Copenhagen, maria.juul.hansen@econ.ku.dk [♯]Australian National University, fedor.iskhakov@anu.edu.au [△]Georgetown University, jr1393@georgetown.edu [†]University of Copenhagen, bertel.schjerning@econ.ku.dk.

The authors would like to acknowledge funding from Realdania, DREAM and the URBAN research project, financed by Innovation Fund Denmark.

1 Introduction

Denmark belongs to a large group of countries that are undergoing a process of strong urbanization and spatial concentration of economic activity. While this has led to increased productivity in the larger cities through agglomeration mechanisms (such as better accessibility of firms to both their markets and supply of specialized labor) it has also resulted in several major societal challenges, including the large and systematic flows of people and jobs with increased traffic congestion and large increases in house prices in urban areas. The result is a changed demographic composition of cities and increased regional inequality. The steady decoupling of urban and rural housing price trends is a clear testament to the latter effect and it has led to an increased inequality in wealth across regions.

A number of policies have been suggested to ameliorate some of the downsides of this development, including infrastructure investments and relocation of government jobs from Copenhagen to the rest of Denmark. However, the dynamic effects of such policies are not well understood due to the complexity of households' commute choices, job and residence mobility, and their interactions with the housing market. Dynamics are crucial for moving and job location decisions since these are made under uncertainty about future house prices and job opportunities, and due to the substantial fixed moving costs that are implied by partly irreversible investments in property and the cost of searching for a new job. These long-run implications of location decisions together with key life events imply that intertemporal incentives are likely to underlie much of the observed behavior. It would be unrealistic to assume that households are not forward-looking and that location choices and housing demand are stationary through the life cycle. Due to commute costs, we must keep track of these mechanisms simultaneously to credibly predict the quantitative consequences of such policies on local house prices, commuting patterns, residential sorting, and inequality. This poses a key challenge that we will address in this paper.

We develop and estimate a dynamic equilibrium life cycle model of residential and work locations to investigate how individuals choose the location of their job and residence, and how urbanization affects house prices, commuting patterns, and demographic composition of cities. We use this model to consider the effects of a number of changes to the economic environment including the effects of an increase in the local housing supply and increased cost of commuting between all residential and work location regions. In doing so, we consider the implications for job mobility, residential sorting, commuting patterns, local housing demand, and house prices.

The modeling framework is a structural life cycle model formulated at the individual level, where people simultaneously choose residential and job locations by taking into account their need for housing, their wage potential at different job markets, commuting costs, amenities, and moving costs. Our model is inspired by the work of Buchinsky et al. (2014), which we extend to a dynamic discrete-continuous choice setting with endogenous house prices and equilibrium constraints, where households make a continuous choice of house size and discrete choice of location of work and residence. Commuting is endogenously determined by the distance between chosen work and residence, and house prices are determined in equilibrium.

By using a life cycle model, we account for individual heterogeneity. To avoid the complexity associated with a complete modeling of life cycle consumption, savings, and borrowing decisions, we assume individuals have quasi-linear utility functions and do not face any borrowing constraints. Instead, we approximate some of these effects by allowing marginal utility of money to be increasing in individual income to reflect that richer households have a higher demand for housing and sort into more expensive regions.

Given the complexity of the model, we also abstract from fully modeling the dynamics of the housing size decision, and we assume that households ignore any adjustment costs of changing house size. The abstraction from such adjustments costs as well as the stylized modeling of consumption, savings, and demand for house size, effectively means we do not distinguish between home owners and renters, but rather model everyone as renters who pay a fixed share of the total house price each period. This essentially amounts to assuming that renting a home or owning a home are nearly equivalent with the "rent" a homeowner pays consisting of the sum of mortgage payments and the opportunity cost of the equity capital the owner has in the home.

The described simplifications allow to keep the model computationally tractable while studying location decisions in more detail. With quadratic utility of house size we can derive a closed form for the optimal amount of housing in each residential location as the solution to a static housing subproblem that can be solved independently of the overall discrete dynamic programming problem governing residential location choices. Given the solution of optimal choice for house size for a given residential location, all dynamic decisions are discrete (job and residential locations) and all of the state variables of the model are also discrete.

Even with these simplifications a challenge of the model presented in this paper is that both the number of states and choices are proportional to the number of combinations of work and residence locations. To ameliorate this curse of dimensionality and avoid solving the model for extremely many combinations of states and choices, we aggregate location choices to the municipal level and restrict attention to the island Zealand (which includes Copenhagen and its soundings). Out of the 98 municipalities in Denmark, we consider the 18 municipalities located in the Greater Copenhagen Area in detail as well as the outside region (rest of Zealand).

We estimate this model using high-quality Danish administrative data that allows us to track the entire population of households, its members, their jobs, and residential locations for the period 1992-2016. In the estimation we focus on the subperiod 2005-2010. These data contain very detailed linked information about location and size of houses, individual employment, wages, and residential and work location dynamics for all individuals and households in Denmark. To estimate the model and compute the equilibrium house prices, our model is repeatedly resolved for many types of individuals as a subroutine of both i) a structural nested backward induction maximum likelihood estimation routine to estimate the preference parameters, and ii) an equilibrium solver that finds paths of housing prices that equate the demand for the available supply of houses in Denmark implied by a microaggregation and simulation of the model based on the parameters estimated in part i).

We assume a fixed supply of housing in the short run and thus abstract from the longer run dynamics where new houses are built in response to changes in house prices. Given that the supply of housing is quite inelastic compared to housing demand, we think this is a reasonable approximation in the shorter run. Using this model, we compute the effects of

i) increased housing stock in the two most urbanized areas of Denmark, and ii) increase in marginal cost (per hour of) commuting between residential and work-location regions.

We find that i) results in increased degree of urbanization as households move from the peripheral regions towards the center. Equilibrium prices fall in all locations, especially in the two locations where the policy was implemented. ii) implies lower commuting on average among employed individuals, but also increases the share of non-employed individuals, especially for those residing in the remote regions, where the equilibrium prices also drop.

The rest of the paper is organized as follows: Section 2 gives an overview of the existing literature and summarizes our contribution relative to existing studies. Section 3 presents the data sources and describes the institutional setting. It also provides descriptive evidence of house prices, residential and work location choices, and the resulting commuting and spatial sorting. Section 4 outlines the model. Section 5 introduces the algorithm that we use to solve and estimate the model. Section 6 describes how we solve for equilibrium prices in the short run. Section 7 presents the parameter estimates and model fit and makes a number of counterfactual simulations focusing on how changes in the local house stock, job density, and commuting costs affect house prices and optimal location decisions. Section 8 concludes and gives directions for future research.

2 Related Literature

This paper builds on and contributes to several strands of the literature covering theoretical and empirical models of location choice in continuous and discrete settings. This section provides a short review of the literature, leading to the dynamic equilibrium model of simultaneous choice of both residence and work location that we develop in this paper.

2.1 Monocentric city model

The literature on household location decisions is based on theory and methodology developed in industrial organization and labor economics. The literature deals with sorting, i.e. with the mechanism that market forces make people with similar preferences and personal characteristics self-select and cluster in certain locations. The urban economics literature is a separate research field dating back to seminal papers by Hicks (1932) and Sjaastad (1962). They made economists interested in understanding the driving forces and implications of how individuals locate. But there were also other contributions that led to this rising interest. Tiebout (1956), was the first to argue that when people sorted ("voted with their feet") in terms of residential location they implicitly revealed their demands for local public goods that were exclusively available in different locations. He focused on the effect that fiscal competition had on income sorting between jurisdictions. At the same time Alonso et al. (1964) developed the monocentric city model, which was enriched by Mills (1967) and Muth (1969). In contrast to the Tibeout model, this was a model for income sorting across geographical space and has become the foundation for many analyses of locations within a city. The main idea was to consider job locations to be exogenous at the center of the city which reduces the residential problem to a choice of how far to commute to one's job, thus ignoring any other travel time and distance that the individual might use to decide on his optimal location (e.g. travel to shops, family or daycare). In a strict sense, the model thereby took as given that people like to live in big cities, but does not explain why they wish to live close to the Central Business District (CBD), except the fact that commute is shorter since all jobs are located in the CBD. It therefore focused on the trade-off between living close to the CBD to get a shorter commute at the cost of more expensive housing. Overall, the consumer maximization problem is standard, except that consumers also choose a residential location (a distance from the CBD) on top of the optimal amount of housing and a composite good. The housing prices at each location respond to offset the marginal decrease in utility that stems from living further away from the CBD. Besides modelling the consumer behaviour, a construction sector which builds houses by use of land and capital is part of the set-up. Land prices are therefore endogenously determined in equilibrium and the model predicts that as land prices increase closer to the CBD, construction firms tend to build with a higher density, i.e. to build tall apartment blocks rather than one story houses.

Extensions of the monocentric city model framework include work by, among others, McMillen (2006) and Ahlfeldt and McMillen (2015) (who focus on the intensity of development) and the Ogawa and Fujita (1980) model that endogenises the location of firms as well to explicitly model agglomeration economics. These other branches of the urban economics literature are out of the scope of our short-run study.

The general scope of the monocentric city model is to predict how geographical space is divided between residential and land use and thus to predict the size of urban areas. The urban area increases until the marginal value of devoting more land to cities equals the marginal value of decreasing agricultural land use. The most prominent recent paper within this branch of the literature is Ahlfeldt et al. (2015). They set up a general equilibrium model of internal city structure where people select a combination of residence and job and where wages and prices of land adjust in response to moving patterns. Their focus is on estimating the extent of agglomeration on productivity, not on the location choice per se. Even though they do study how equilibrium land prices change in response to altered moving patterns, they can only estimate the long-run effects in a static modeling framework.

2.2 Reduced-form models

The discrete choice framework that has been increasingly popular in recent decades was initiated by McFadden (1974). He provided a methodology for analyzing choice behaviour when the agent optimizes with respect to choices from a discrete rather than continuous set as in the monocentric city approach. He was also the first to really contribute to the discrete sorting literature in McFadden (1978), where individuals choose a specific location rather than just a certain distance from the CBD.

Another distinction within the location choice literature is that it originally centered around two types of models: human capital on the one hand and hedonic models on the other. Human capital models are, among others, motivated by Topel (1986). Hedonic models were introduced around the same time as McFadden came up with the discrete choice estimation methods, namely by Rosen (1974) and further extended by Roback (1982). Rosen (1974) set forth a method for estimating marginal willingness to pay (MWTP) for goods for which there were no formal markets such as air pollution, crime rates and scenic views. Introducing hedonic price functions, he explained how researchers could use data on observed location choices by households and housing prices for the different locations to compute implicit price indices of these non-traded amenities. Whereas the human capital models argue that migration occurs due to disequilibrium in the labor market (people move to a new location to earn a higher wage), the hedonic approach asserts that individuals might move even if housing and labor markets are in equilibrium because they might have changed demand for location-specific non-traded amenities. While the human capital literature sees earnings differentials as temporary circumstances that will mitigate when workers relocate in order to get the highest possible return on their human capital investments, hedonic models explain how wage and housing price differentials may not be completely eliminated, as they may fail to compensate individuals for location-specific (dis)amenities. With the emergence of the sorting literature, which has been thorougly surveyed by Kuminoff et al. (2013), these two approaches were combined into a unified framework.

2.3 Structural static models

To explicitly take this sorting into account, a growing number of papers structurally model location decisions, though until very recently mainly by using static models. One example is Borjas (2000) who looked at how immigrants affect the equilibrium in the local labor markets across several geographic areas. This paper points out that the possibility of moving to another location for work among natives is not sufficient to cancel all wage differentials across locations. This is because people face high moving costs that make them reluctant to move for the best wage offer. In contrast, immigrants from other countries do not to the same extent incur moving costs on top of those associated with leaving their home country and are therefore more prone to settle and work in the area characterized by the best wage offer. Bayer et al. (2009) is another example which stresses the importance of not only relying on the first-order conditions from the traditional hedonic model, but rather combine the frameworks originally presented in McFadden (1978) and Rosen (1974). This allows for explicitly accounting for moving costs which is necessary in order to get unbiased estimates of MWTP for non-traded amenities. Besides the change in amenities, the rising prices in cities across the world has also led researchers to study the effect of altering the housing supply, including Nathanson (2019) which estimates a static equilibrium model of residential location, while the job search behaviour across local labor markets has been studied in Manning and Petrongolo (2017). They find that there is a sharp decay in attractiveness of jobs as they get further away from the home, which also speaks in favor of modelling the decisions of home and job jointly as we do in this paper.

However, the combined modelling of work and home locations has not been the focus of the literature until recently. Tsivanidis (2019) is a very recent exception. He estimates a static general equilibrium model of home and work locations as well as car ownership and housing size to quantify the effects on sorting of workers and their welfare from a large infrastructure investment. He documents that it is indeed essential to account for the spatial reallocation of workers and general equilibrium effects as we do in the model of this paper. A related question is discussed in Albouy and Stuart (2019) which decomposes the determinants of residential location choices into a number of amenities using a neo-classical spatial equilibrium framework. They conclude that quality of life (i.e. factors related to the utility of residing in a region) are more important than trade productivity (i.e. determinants that affect people's taste for jobs there), but that both have an effect.

2.4 Structural dynamic models

The lack of appropriate data and computational difficulties are the main reasons why the literature has focused on static models for so many years. Dynamics are crucial, however, as outlined in Section 1. Kennan and Walker (2011) were the first to add dynamics to a structural model where individuals optimize over a set of residential locations each period. There were a few predecessors in the dynamic location choice literature such as Holt (1997) and Tunali (2000). However, they both did not distinguish between alternative locations, but modeled only the move-stay decision. Dahl (2002) did do so by allowing individuals to choose between all U.S. states, but individuals only made one moving decision for their entire life. Gallin (2004) looked into how changes in expected future wages affected net migration in an area, but he used aggregates and thus did not model how the individual responded to this. Lastly, Gould (2007) studied how workers choose between residential locations and occupations, but only distinguishes between rural and urban locations.

Kennan and Walker (2011) was the first paper to broaden the application to a more detailed setting, where they allow for many different locations (U.S. states). However, they restrict people to live and work in the same location. They find that better income prospects associated with moving to another location is an important driver of migration decisions. Bishop (2012) also uses a dynamic model but has another focus: namely to set up an equilibrium model to estimate willingness to pay for air quality while controlling for moving costs and forward-looking behavior. In her model, individuals are forward-looking with respect to local amenities, and they can move in expectation of how they evolve over time. Her model can also result in a huge choice set, but she does not address the question of how work and residential locations are interconnected. Bayer et al. (2016) adopt her approach, but go one step further and estimate the willingness to pay for several non-traded amenities of a neighborhood. Additionally, their paper allows for household wealth to evolve endogenously with housing prices such that households' expectations about future housing prices can affect their decisions. Location choice hence becomes dynamic both due to moving costs and wealth accumulation.

Another recent contribution to the literature is Oswald (2019). He models the choice of consumption, home ownership, and residential location. Whereas there has been a tradition of ignoring the choice of home ownership, he integrates this decision into the model to account for the fact that home owners' wealth declines when house prices do, while renters may benefit from lower rents. He argues this is important as the option of moving is a way to self-insure against local shocks to the housing and labor market and estimates the value of this self-insurance mechanism. The question of owning or renting is also taken up in Favilukis et al. (2019) which calibrates a rich dynamic equilibrium overlapping generations model for commuting, consumption, housing, residential location, and own/rent decisions for New York City. They use the model to assess the implications of zoning and rent control policies. However, they assume away moving costs and assume a two-alternative choice set for residential locations. Halket et al. (2015) studied the allocation of properties to ownership and rental from a supply side perspective, while Attanasio et al. (2012) added a choice of housing size and consumption over the life cycle on top of the choice of owning or rentning. They abstract from the discrete location decisions though, but find that demand for housing does indeed react to prices and income shocks.

While the papers mentioned above do not model detailed labor markets, a number of papers concentrate on the dynamic aspects of migration decisions, employment status, and job search such as Ransom (2016), Schmutz and Sidibé (2015) and Mangum (2015). The latter two include equilibrium constraints on the labor market and real estate market, respectively, but none of them model the joint decisions of home and work. Guglielminetti et al. (2017) model the job search process for unemployed people and take home location into account, but for two possible locations.

Buchinsky et al. (2014) are the first to structurally estimate a dynamic model of residential location that also adds the choice of work location. The model mainly builds on Kennan and Walker (2011), but a very important extension is that they distinguish between home and work locations. Hence, individuals choose home and job locations as well as labor market status and sector each period. By relaxing the assumption of zero commute they were able to model commute costs. Moreover, they allow people to have expectations about the job offerings before they decide on their locations. This complicates the empirical implementation of the model since job offers are not observed, but also adds a more realistic aspect to the model. We employ a similar approach in the model of this paper. Even though their model is very rich in terms of its choice set, it is a partial equilibrium model where both wages and prices are taken as exogenous. The authors provide arguments why this is not too important for their setting. Another limitation of the paper is that it restricts attention to the case of hihgly educated immigrants from the Soviet Union migrating to Israel, which is likely to be a very selected group. Their results therefore mainly regard immigrants instead of people who migrate within a country, which is the focus of the current paper.

3 Data and Institutional Background

This section provides the description of the data we use to estimate the model. We also provide descriptive statistics on property prices, urbanization, overall life cycle patterns of moving home and job and demand for housing size, while we use Section 7 to go into more details of the location and sorting patterns when we present the model fit from the estimation.

3.1 Danish register data

We use administrative data provided by Statistics Denmark which holds information on every individual living in Denmark in the period 1998-2010[1], although for estimation we use the years between 2005 and 2010. The personal registers contain information on individual background characteristics like home and work address, education and number of kids, while other registers hold data on home sales and prices, home owners and dwelling characteristics. Since we focus on the Greater Copenhagen Area in the current paper, we redefine regions according to Figure 1. We include everyone who has either work or home region within the Greater Copenhagen Area[2]. Below we go over each of the separate data sources.

Figure 1: Definition of regions

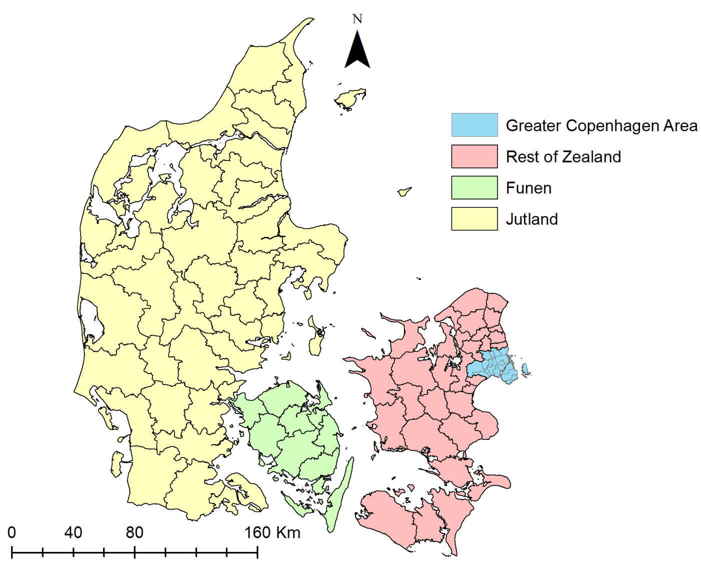

[1] But can be extended to 1992-2016.

[2] Appendix A provides more details on the geographic units in Denmark. Including the entire Denmark is left for future research.

In 2007, a municipality reform took place in Denmark which reduced the number of municipalities from 271 to 98. We are able to track how municipalities were combined, and use the more coarse definition for all years.

Register data of individual background characteristics is recorded on January 1st each year and list all individuals who are officially registered with an address in Denmark. Each individual in the registers is associated with a family identifier[3] and an anonymized version of their official social security number which allows linking of different registers on the individual level. We use the age, gender, address identifier[4], whether she has children, how old the children are, and if she lives with a partner (we track both marriages and cohabitations) as background characteristics. The data on workplace and other workplace-related variables such as industry and occupation, along with the wage, are recorded in the end of November, and are linked to the individual data from the previous year. For each individual who has more than one wage-earning job, we use information from the main occupation which is determined by the largest source of income. Individuals who do not work are either classified as unemployed or outside the labor force.

Commute time data come from The Danish Traffic Model (LTM) which has been developed by researchers at The Technical University of Denmark (DTU). They divide Denmark into 907 traffic zones (LTM zones) and modelled commute time between each pair of regions. Since our model is formulated in terms of municipalties, not LTM zones, we compute a commute time measure by each transport mode between any pair of LTM zones within a municipality pair. For a given pair of LTM zones in a municipality pair, we use the commute time from the mode with the shortest commute time. We then weigh the commute time of each observed LTM pair in the municipality pair with its estimated number of trips by that mode from the traffic model and thereby get a trip-weigted average commute time between any pair of municipalities.

Information on property prices come from the sales transaction register. We deflate all sales prices by the consumer price index with 2011 as the reference year. We use only private

[3]Families are defined as everyone on an address who are related biologically, registered partners, or opposite gender couples. Singles are families of one. However, for each individual we also observe the identifier of their partner.

[4]Addresses are anonymized within a region, but associated with unique identifiers.

sales and disregard properties with commercial-only purpose. A more detailed description of the Danish register data is available in Appendix F.

3.2 Property prices

Property prices in the Copenhagen area, and especially in the center and northern parts, have tended to be higher than prices in the remaining regions of Denmark for the last three decades. Since the begininning of the 1990s, however, the prices by regions have diverged. The hierarchy of regions from lowest to highest price per square meter is more or less unchanged over time. But the prices in central parts and north of Copenhagen and its nearest surroundings started to increase in the mid 1990s, the rest of the country experienced much more modest or stagnating prices. Figure 2a documents this evolution since 1998 to 2010. In the years just before the financial crisis, prices of all regions began to rise, but still steeper for areas close by Copenhagen. The gap between prices of the most urbanized areas compared to the suburbs and rural areas widened as a result. While prices in the center of the GCA (Copenhagen and Frederiksberg) continued to grow until the outset of the financial crisis in 2006, the flow of people to that same area showed the reverse trend until 2005, cf. Figure 3. As prices reached the highest level in decades in 2005, the net outflow from the center of the GCA topped at 3,000 people. From 2006, on the other hand, the net outflow rose towards 0 in 2010 while prices dropped significantly during the years 2006-2009. Historically, we have therefore seen increasing prices both in times of increasingly negative net inflows to the center of the GCA and in years where the negative net inflow was getting more modest.

Another stylized fact about prices is how they relate to the size of the homes. As Figure 2b depicts, prices are almost a linear function of square meter living space. However, the strength of the relationship differs across regions with the Copenhagen, Frederiksberg and Gentofte north of Copenhagen showing the steepest relationships. On Zealand, which is a much less dense area, the slope is almost five times lower than in Copenhagen.

Figure 2: Prices by year and square meters for selected regions

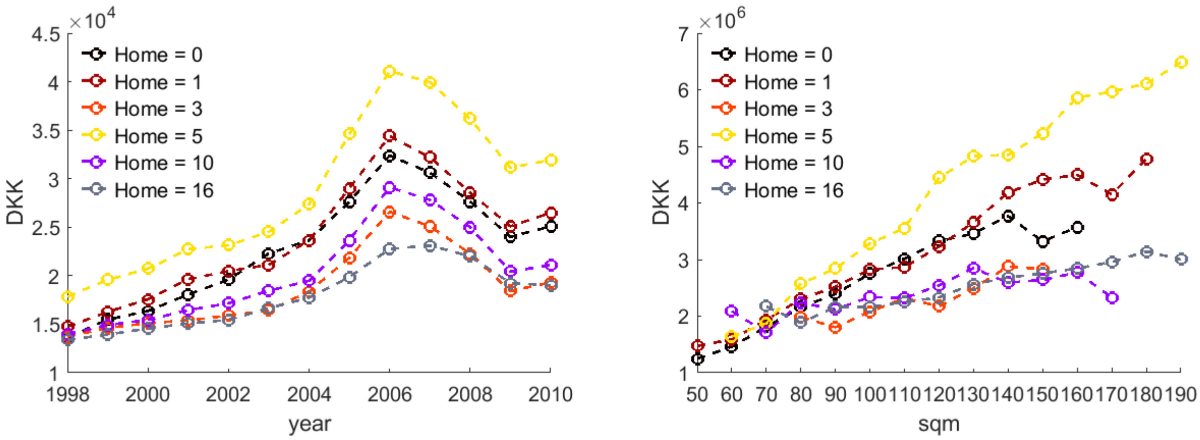

(a) Price per sqm

(b) Total price by sqm

Note: $Home = 0$: Copenhagen municipality, $Home = 1$, Frederiksberg, $Home = 3$: Broendby, $Home = 5$: Gentofte, $Home = 10$: Hvidovre, $Home = 16$: Rest of Zealand. The figure shows real sales prices deflated by the 2011 consumer price index.

Figure 3: Sales price per sqm and net in-migrants for Copenhagen and Frederiksberg regions over time

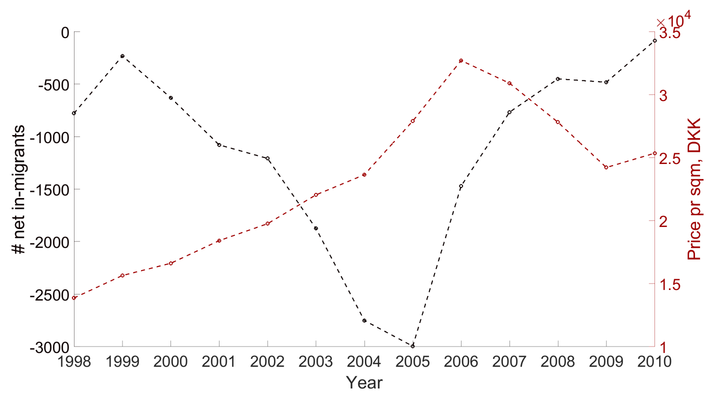

Note: Real sales prices deflated by the 2011 consumer price index.

3.3 Life cycle patterns

In this section we summarize descriptive statistics on home and work location choices, commuting and house size demand. We show a very clear life cycle profile on all margins. This underlines the need for modelling these decisions as being affected by dynamic incentives.

Figure 4: Share moving home region by age

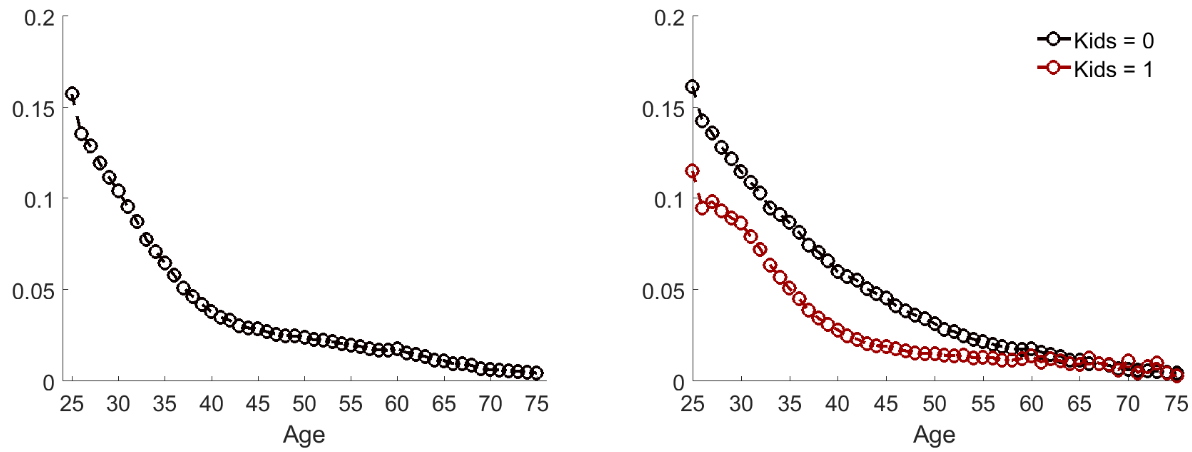

(a) Aggregate (b) By children

3.3.1 Home decision

On average over the life cycle, 4.3 percent of people move to another home during a year. These moves only include moves across municipality boundaries. Intra-regional moves are thus disregarded, but both renters and home owners are included. The moving probability changes a lot over the life cycle as Figure 4a shows. While 25-year-olds have a 16 percent probability of moving to another home region, this drops almost linearly until the late 30s where the probability is about 4 percent. These numbers also depend significantly on the parental status of individuals. Those with children have a steeper decline in moving probability from the beginning of their 30s until the mid-40s compared to those without children. At age 35, 5 percent of people with kids move their home to another region compared to 9 percent of those with no children. By the end of the 50s the shares are almost identical across the two groups, consistent with the fact that most individuals no longer have children living at home anymore. The clear life cycle perspective in moving behaviour speaks in favor of using a dynamic model.

There is furthermore a clear life-cycle profile in sorting patterns as we illustrate in Figure 5a. It shows the share of individuals living in Copenhagen municipality and Frederiksberg by age and schooling level. Clearly, young and highly-educated people are particularly more likely to live in these municipalities which make up the centre of the GCA. One attraction is that

Figure 5: Sorting into GCA centre, Copenhagen and Frederiksberg

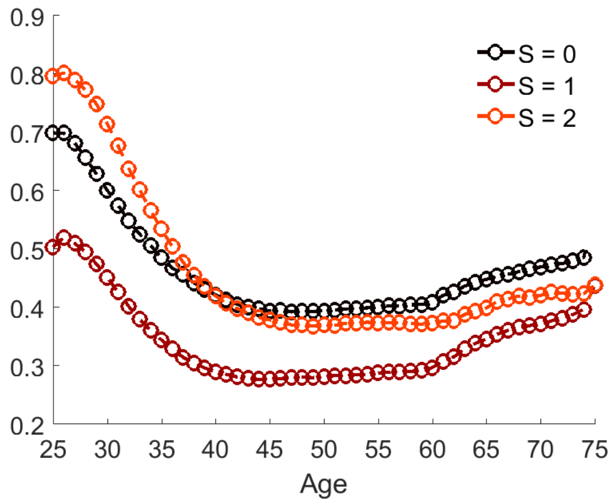

(a) Share of GCA population living in Copenhagen and Frederiksberg

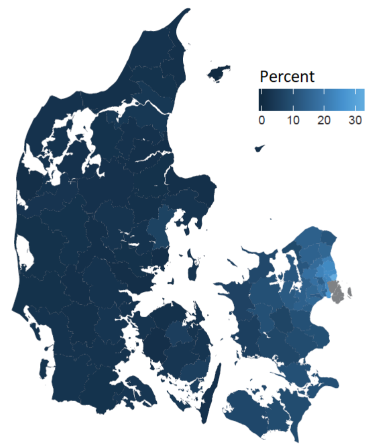

(b) Inflow by home $t-1$

Note: In (a): $S=0$: low education, $S=1$: medium education, $S=2$: high education. In (b): shows the distribution of in-migrants to Copenhagen and Frederiksberg across home regions in $t-1$.

the GCA educational institutions and universities are primarily placed there, but it is also the place with most high-skilled jobs. This can be seen in Table B1 in appendix which shows the average job density by work region and education group. At the age of 25, 80 percent of the high-skilled people live in Copenhagen and Frederiksberg compared to 70 percent of medium-skilled and 50 percent of low-skilled individuals. The probability stagnates at around 40 percent for high- and medium-skilled by age 40 and at 30 percent for low-skilled people.

Turning to Figure 5b, it shows the distribution of home locations at time $t-1$ for people who move to Copenhagen and Frederiksberg at time t. Clearly, the main part of in-migrants come from municipalities located close by Copenhagen and Frederiksberg. The probability of originating from a region is thus decreasing in the distance to Copenhagen and Frederiksberg. Odense and Aarhus stand out as they comprise a relatively large share of the locations from which the in-migrants originate given their distance from Copenhagen. This indicates that people from other big cities of Denmark are attracted to Copenhagen despite the distance, but overall less than 22 percent of the in-migrants come from Funen or Jutland. This underlines that when we focus on Copenhagen in the estimation it is less important to model location

Figure 6: Share moving work region and average commute time

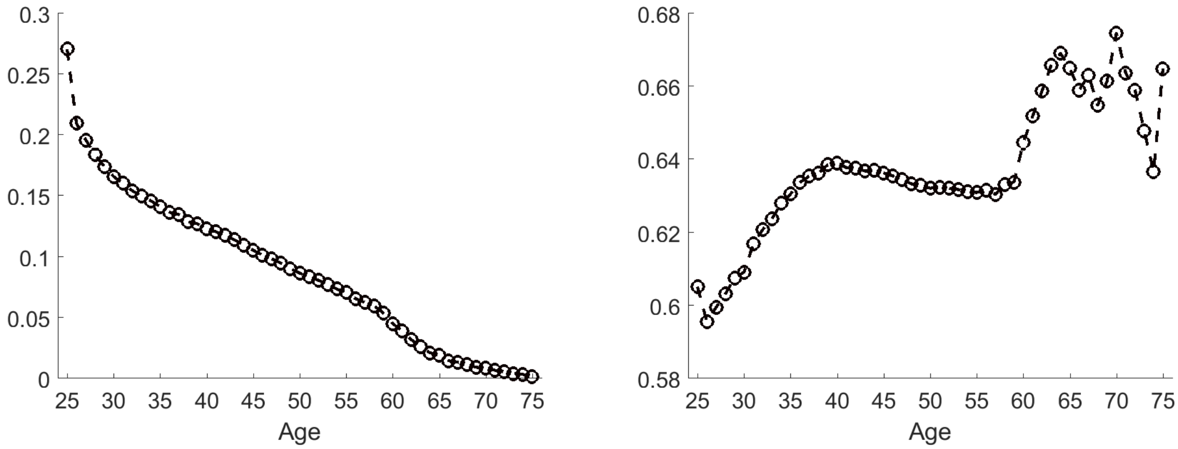

(a) Move work region

(b) Commute time (hours)

decisions for people living on Funen or in Jutland in detail.

3.3.2 Work decision

The average share of people changing job locations from one municipality to another is 14.5 percent. As for the home moves, this share changes over the life course as Figure 6a pictures. While the 25-year-olds in the data have a 27 percent probability of moving to a job in another region (excluding transitions in and out of unemployment), the probability is 16 percent at age 30. From then it falls linearly to 6 percent at age 59 and then drops towards zero from there. The sharper drop from age 59 to 60 is due to people being eligible for early retirement benefits at that time.

By linking home and work locations we get the commute time. As Figure 6b shows, the average commute time is increasing from age 26 to 40 whereafter it starts to decline slightly. There is a sharp increase in commute time at age 60 until 65. This is explained by self-selection of workers who stay on the labor market even after they are eligible for retirement.

Having covered residential and work locations separately, Figure 7 shows the share of residential moves which are associated with a job move. For 58 percent of the inter-regional moves, there is no change in work region. However, 33 percent do change job either the year before, in the same year or the year after they relocate their residence. The remaining 9 percent change jobs more than once during that time window. This finding underlines the importance of modelling residential and work location decisions as joint decisions.

Figure 7: Probability of moving home region at t by number job moves in $t-1, t$ and $t+1$

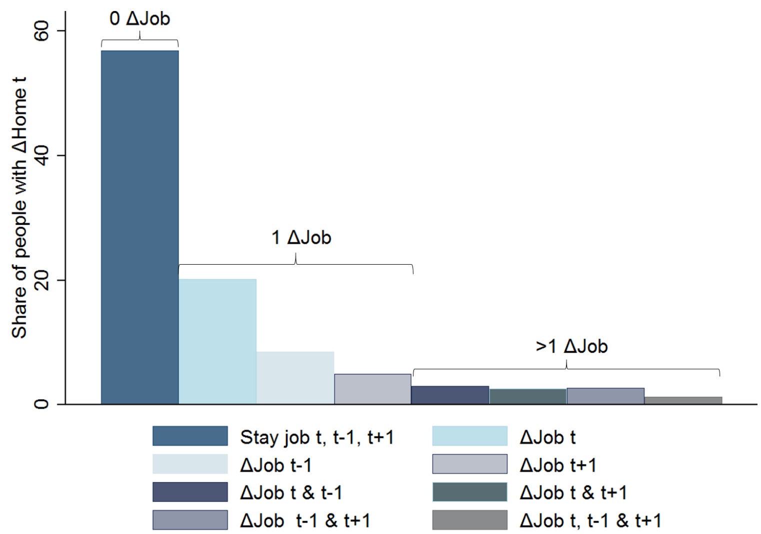

Note: ΔJob refers to "change in work region". t is the time of the home region move.

3.3.3 Housing demand

Another important aspect of home regions, besides square meter prices, is house sizes. Figure 8a illustrates how the average square meter demand evolves over the life cycle: it starts at 75 square meters at age 25, increases to 115 square meters in the beginning of the 40s, then levels off until the late 50s where it begins to decline. Multiple-member households generally demand more square meters throughout their lives as illustrated in Figure 8b, which distinguishes between individuals with and without children, and Figure 8c, which separates demand by marital status. Turning to Figure 8d, there clearly is a gradient in income too: higher-income people demand more square meters also when conditioning on their home region. People in Rest of Zealand therefore live in bigger houses for any given income bracket compared to those who live in e.g. Copenhagen. This is closely related to the spatial variation in house prices documented above; regions with a high square meter price are characterized by smaller houses, all else equal. Due to the very differentiated demand for square meters across regions, it is important to explicitly model this when modelling the home location choices.

Figure 8: Housing size demand by age and income

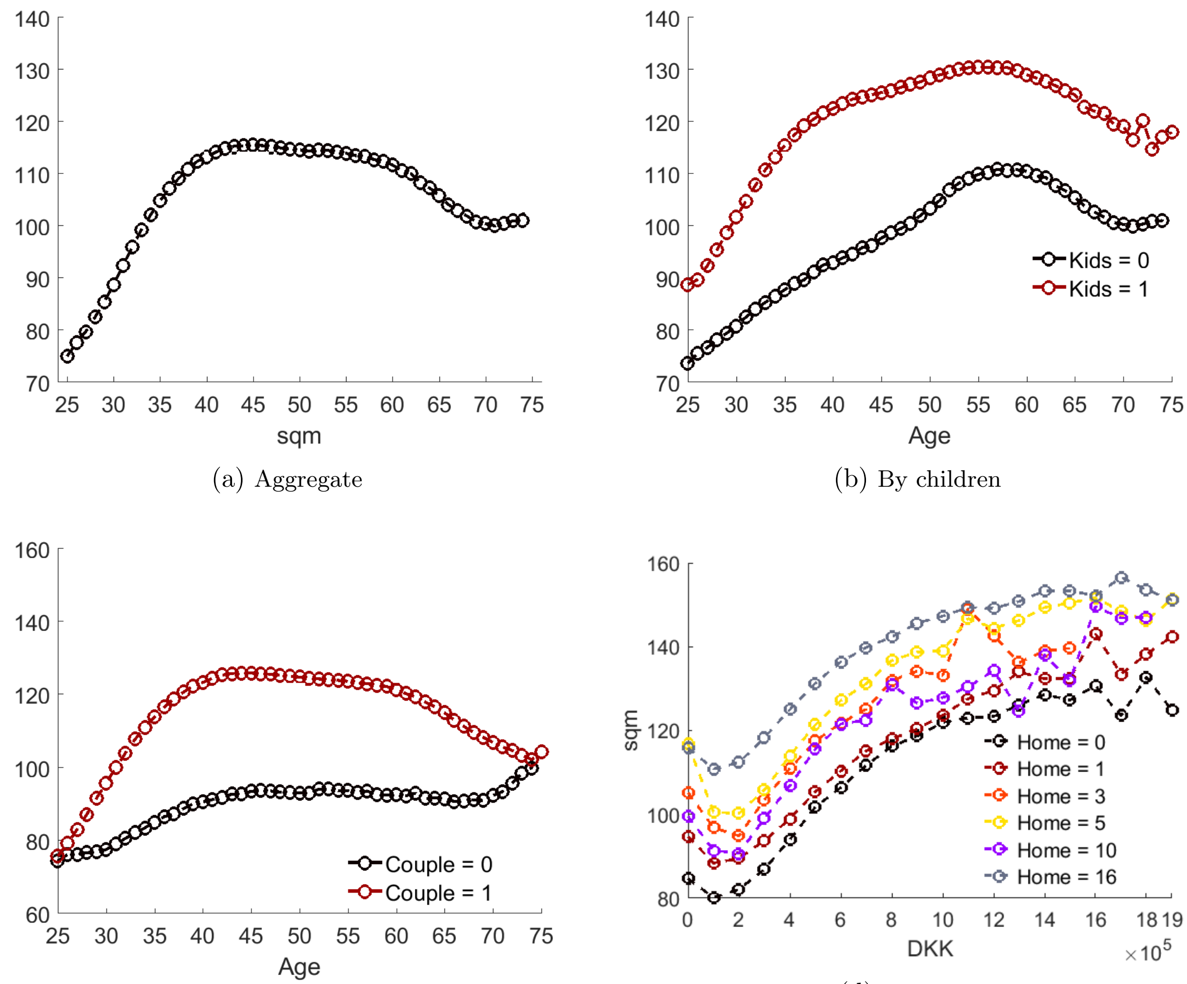

(a) Aggregate

(b) By children

(c) By marital status

(d) By income

Note: $Home = 0$: Copenhagen, $Home = 1$, Frederiksberg, $Home = 3$: Broendby, $Home = 5$: Gentofte, $Home = 10$: Hvidovre, $Home = 16$: Rest of Zealand. The figure shows real sales prices deflated by the 2011 consumer price index.

4 A Dynamic Model of Residential and Work Locations

In this section we lay out the model of individual housing demand and location choice of work and residence, formulated *conditional* on prices and job opportunities in different regions. The individual choice model is embedded into the general equilibrium model in Section 6, where we derive the housing market equilibrium based on the solution to the individual problem.

4.1 Sequential choice of work and residence locations

Consider an individual decision maker who in each time t of her lifecycle, $t \in \{t_0, \ldots, T\}$, chooses work and residence locations, as well as the size of residence. The decision maker in our model has unitary preferences, and we interpret t as the individual's age. Time horizon $T = 75$ is chosen to be sufficiently large such that very few changes of resident or work location occur after this age.

Modeling location choice is computationally burdensome when the number of regions is large. Denoting R the number of regions, even without the choice of house size, the number of discrete location choices for residence and work is R^2. Moreover, because the cost of moving is an unavoidable element of any realistic location choice model, the state space of the model necessarily includes the previously chosen locations, thus requiring R^2 states before any of additional heterogeneity is accounted for. Therefore, in order to keep the model tractable, we introduce a number of simplifying assumptions right from the outset.

First, we assume that there are no fixed costs associated with scaling the size of a house up or down. The decision maker only pays the rental value of a home, calculated by its size times local square meter prices, and she is in each period free to resize her house without moving. It follows that the first order conditions fully characterize the choice of the house size conditional on the individual's characteristics and the attributes of the region of residence, and it can thus be expressed analytically and substituted into the indirect utility function. We derive the corresponding static demand for home size in Section 5.2 below. In presenting the dynamic discrete choice set-up in this section we therefore abstract from the continuous choice of house size, which is subsumed by a general indirect utility specification.

Second, we assume that work and residence location choices are made sequentially, namely that work location choice is made first followed by the residence location choice. Even though this assumption does not immediately decrease the number of alternatives in the resulting nested choice model (R nests by R alternatives each), it allows us to introduce a sensible job matching process. Namely, we differentiate between the job transition *choice* that denotes intention, from the job *outcome* that becomes the next period work location. In other words, our model allows for unsuccessful attempts to change work location and involuntary unemployment.

In our computational approach we recognize the fact that the expected future value of the current period choices only depends on the work and residence locations *realized by the end of the period*. We therefore formulate the dynamic programming problem in terms of expected value functions, keeping its dimensionality on the order of R^2 rather than R^4 (R^2 states by R^2 choices) as would be required by the traditional solution in the space of choice-specific value functions.

Because "work location" appears in the model in three different forms (existing, intended, and realized work location), we use the following explicit notation to distinguish between them. We denote wl_t the beginning of the period existing work location (*state*) and d_t^w the period t choice of intended work location (*choice*). This may or may not be the same as wl_t. Finally, to denote the *outcome* of the job match process during period t we simply use wl_{t+1}, as the realized in period t work location becomes the existing one in period $t+1$.[5]

To maintain uniform notation, rl_t denotes the existing period t residence location and, correspondingly, d_t^r denotes the choice of new residence location. We assume perfect control over the location of the residence (subject to the equilibrium house prices), and therefore the location of residence in period $t+1$ is given by the choice at period t, i.e. $rl_{t+1} = d_t^r$.

The precise timing convention we use as follows. Each period t starts off with a given work and residence location, and other variables x_t to be described below, forming the vector of state variables $s_t = (wl_t, rl_t, x_t)$. We assume that individuals make their work and residential location choices sequentially but instantaneously at the start of each period t, with the intended work location decision made first, followed by the residential location decision made *conditional* on the realization of the employment search, i.e. realized work location wl_{t+1}. Once the intended work location is chosen, the job search outcome is realized, and the

[5] Using notation wl_{t+1} as the realized work location in period t involves a degree of confusion with the time subscripts, but we opt to bear this cost to avoid having an additional outcome variable.

residence location choice is made, the household determines the optimal house size depending on their own characteristics and the chosen region of residence. Thereafter, the housing consumption is enjoyed for the rest of the period, and the process transitions to the next period. We describe the transition rules and list all the components of the state vector s_t in Section 4.4 below.

Following the tradition of the discrete choice literature, we assume that the choices of both work and residence locations depend on the *IID* extreme value idiosyncratic shocks $\epsilon_t = (\epsilon_t^w, \epsilon_t^r)$ that can be interpreted as the components of the utility that the econometrician does not observe. We assume that these stochastic components are revealed to the individual sequentially: at the time of the work location decision d_t^w only the "work location shocks" ϵ_t^w are known, whereas the "residential location shocks" ϵ_t^r are only revealed after the individual learns the outcome of their employment search. In other words we assume that the households find out the idiosyncratic attributes of the residence locations only once they know where their job takes them. These assumptions lead to the standard nested choice structure of the work and residence location decisions, with standard analytic expressions for choice probabilities at each level, and inclusive values of the residential choice given by the McFadden's social surplus (*logsum*) functions, as we discuss below.

4.2 Specification of the job search process

Before deriving the recursive formulation of the model, we specify the possible transitions in the job search process. A spatial model with fixed wages could lead to the outcome where far more people want to move into a high wage region than there are available jobs. Introduction of the labor market into the model allows us to avoid this unrealistic scenario.

Let the spatial work region $wl_t = \emptyset$ denote the state of non-employment, which can naturally be combined with any residence region rl_t. We assume that unemployment can be chosen voluntarily, but also allow for involuntary job separations with a certain probability, including the cases when no job transition is intended ($d_t^w = wl_t$).

Let $\pi_t^n(d_t^w, wl_t, x_t)$, $d_t^w \neq wl_t$, denote the probability of successfully finding a new job in the region d_t^w, conditional on the household characteristics and other variables in state vector

x_t. For the case $d_t^w = wl_t$, $\pi_t^n(wl_t, wl_t, x_t) = \pi_t^k(wl_t, x_t)$ simply denotes the probability of keeping the existing job in location wl_t[6]. With the complementary probability $1 - \pi_t^k(wl_t, x_t)$ a transition to non-employment $wl_t = \emptyset$ occurs.

If the individual chooses to stop working, $d_t^w = \emptyset$, then $\pi_t^n(\emptyset, wl_t, x_t) = 1$, i.e. there is "perfect control" over the decision to stop working. However, for an individual who is searching for a new job in the region $d_t^w \neq wl_t$ the work location transition probabilities are given by

$$wl_{t+1} = \begin{cases} d_t^w & \text{with probability } \pi_t^n(d_t^w, wl_t, x_t), \\ wl_t & \text{with probability } \big(1 - \pi_t^n(d_t^w, wl_t, x_t)\big)\pi_t^k(wl_t, x_t), \\ \emptyset & \text{with probability } \big(1 - \pi_t^n(d_t^w, wl_t, x_t)\big)\big(1 - \pi_t^k(wl_t, x_t)\big). \end{cases} \quad (1)$$

In words, the transition probability in equation (1) says that if an individual chooses to search for a job in some new location $d_t^w \neq wl_t$, then there are three possible outcomes: i) the individual could be successful and receive a job offer in this location; ii) the individual does not get a job offer in the new location but is able to keep her existing job; or iii) the individual's job search is unsuccessful and she is laid off from her current job.

Note the "on the job search" assumption which means that if a worker applies for a job in a different location and does not get that job, they still have the option of staying in their current job, unless they are laid off. The alternative assumption would be to assume that if a worker applies for a job in some other work location, they have to quit their current job first. We think the on the job search assumption is a better approximation to reality: the assumption that workers must first quit their jobs (i.e. have no recourse of staying at their current job if they apply for another job and are unsuccessful) would likely make it artificially risky to change job locations, and such an assumption may result in underprediction of job mobility across different regions.

If the individual does not search for a job in a new location, $d_t^w = wl_t$, we assume that they will continue working in the same location as before, and set $\pi_t^n(d_t^w, wl_t, x_t) = 0$. In this

[6]We abstract from job transitions within the same region.

case (1) takes the form

$$wl_{t+1} = \begin{cases} wl_t & \text{with probability } \pi_t^k(wl_t, x_t), \\ \emptyset & \text{with probability } 1 - \pi_t^k(wl_t, x_t). \end{cases} \tag{2}$$

Specification (1) can be applied for unemployed individuals as well, in which case the last two rows collapse into one, and it takes the form

$$wl_{t+1} = \begin{cases} d_t^w & \text{with probability } \pi_t^n(d_t^w, \emptyset, x_t), \\ \emptyset & \text{with probability } 1 - \pi_t^n(d_t^w, \emptyset, x_t). \end{cases} \tag{3}$$

This specification allows for the possibility that there may be a lower chance of getting a job if a person is currently unemployed, compared to an individual who is currently employed in this location or in some other location. That is, one possible ordering of the employment transition probabilities that might be supported empirically is

$$\pi_t^n(d_t^w, \emptyset, x_t) < \pi_t^n(d_t^w, wl_t, x_t) < \pi_t^k(wl_t, x_t) < \pi_t^n(\emptyset, wl_t, x_t) = 1, \tag{4}$$

so an individual who chooses to stay in their current work location has the highest probability of being employed in this location, $\pi_t^k(wl_t, x_t)$, apart from choosing to stop working $\pi_t^n(\emptyset, wl_t, x_t) = 1$. An individual who is applying for jobs in the region from the outside ($d_t^w \neq wl_t$), has a lower employment probability, while the lowest probability of employment corresponds to the unemployed individuals, $\pi_t^n(d_t^w, \emptyset, x_t)$.

Let $\pi_t(d_t^w, wl_t, x_t, wl_{t+1})$ denote the probability of transition from work in region wl_t to work in region wl_{t+1} (including non-employment $wl_{t+1} = \emptyset$), which encompasses all the transition probabilities described in this section in equations (1)-(3). We present the functional forms assumptions of $\pi_t^n(d_t^w, wl_t, x_t)$ and $\pi_t^k(wl_t, x_t)$ in Section 4.5.

4.3 Recursive formulation and Bellman equations

At each period $t = t_0, \ldots, T$ the individuals in the model maximize the expected discounted utility over the remainder of their life by making sequential work and home location decisions, as well as choosing the house of the optimal size. For every t the attainable maximum is given by the value function $V_t(s_t, \epsilon_t)$ which is a function of state variables $s_t = (wl_t, rl_t, x_t)$ and the stochastic taste shocks ϵ_t. As mentioned in Section 4.1, instead of the value function $V_t(s_t, \epsilon_t)$ we focus on the *expected value function* $EV_t(wl_{t+1}, rl_{t+1}, x_t)$, as a function of the realized work and home locations (after all relocations have been completed). Note that the expected value function at period t depends on the work and residence locations at period $t+1$. Even though this may appear as a type of "clairvoyance" of the decision makers, it is merely the consequence of our timing assumptions. The "next period" location (wl_{t+1}, rl_{t+1}) in fact just denotes the location outcome after the decisions and relocation stage is completed in the beginning of the period. According to the timing convention, during period t, the individual lives at location rl_{t+1} and works at location wl_{t+1}.

Unlike the expected value function $EV_t(wl_{t+1}, rl_{t+1}, x_t)$, the period t (deterministic) flow utility has to account for the switching costs of relocations, and therefore has to depend on both initial locations and the realized location. To allow for a flexible way that switching costs enter the model (both for changing the home and work locations and differentiated for heterogeneous households) in this section we use the generic form of the utility function given by $u(wl_t, rl_t, wl_{t+1}, rl_{t+1}, x_t)$[7]. Note that the choice variables enter into the utility function indirectly: choice of work location d_t^w governs the job search process described in previous section, and under assumed perfect control the choice of residence location, we have $d_t^r = rl_{t+1}$.

Given the nested discrete choice structure in the model described in Section 4.1, the extreme value shocks $\epsilon_t = (\epsilon_t^w, \epsilon_t^r)$ enter the Bellman equation in a non-trivial way. We build the Bellman equation in stages following the backward induction over the events within the time period.

Let β denote the discount factor of the individual. For simplicity we assume it is

[7] Additional assumptions on the utility function could have drastically reduced the computational burden of the model. For example, assuming that the moving costs for residence (c^r) and work (c^w) are additively separable and only depend on the destination, i.e. $u(wl_t, rl_t, wl_{t+1}, rl_{t+1}, x_t) = u'(wl_{t+1}, rl_{t+1}, x_t) - c^w(wl_{t+1}, x_t) - c^r(rl_{t+1}, x_t)$, leads to much simplified expressions of the value functions that can be expressed with the values of households that do not change locations, modified by a collection of constant moving costs, and thus drastically simplifying the computation of the log-sum function and choice probabilities when solving the dynamic programs. However, in order to be able to match the data we have, we prefer to keep the model specification flexible at this stage.

independent of individual survival rates. I.e. we do not take into account that the discounting of future expected values may lower as the individual ages.

Recall that $\epsilon_t^r \in \mathbb{R}^N$ are the stochastic components of the utility corresponding to the choice of residence location, once the outcome of the job search process is revealed, and the new work location wl_{t+1} is known. Let $\epsilon_t^r(d_t^r)$ be the idiosyncratic utility costs/benefits of choosing to move to location d_t^r. We assume it is extreme value with scale parameter σ_r. Let $EV_t^r(wl_t, rl_t, wl_{t+1}, x_t)$ be the *ex ante* expected value for an individual who knows her employment location outcome wl_{t+1} but has not learned the residential location shocks $\{\epsilon_t^r(d_t^r)\}$ yet. This is given by the usual log-sum formula

$$\begin{aligned} &EV_t^r(wl_t, rl_t, wl_{t+1}, x_t) = \\ &\sigma_r \log\left(\sum_{d^r} \exp\{[u(wl_t, rl_t, wl_{t+1}, d^r, x_t) + \beta EV_t(wl_{t+1}, d^r, x_t)]/\sigma_r\}\right). \end{aligned} \tag{5}$$

The implied residence location choice probabilities are given by the multinomial logit formulas

$$\begin{aligned} P_t^r(d_t^r|wl_t,&rl_t, wl_{t+1}, x_t) = \\ &\frac{\exp\{[u(wl_t, rl_t, wl_{t+1}, d_t^r, x_t) + \beta EV_t(wl_{t+1}, d_t^r, x_t)]/\sigma_r\}}{\sum_{d^r} \exp\{[u_t(wl_t, rl_t, wl_{t+1}, d^r, x_t) + \beta EV_t(wl_{t+1}, d^r, x_t)]/\sigma_r\}}. \end{aligned} \tag{6}$$

Now consider the choice of the work location at the beginning of period t, d_t^w. Because this choice is moderated by the job search process, we have to take into account the probabilities $\pi_t(d_t^w, wl_t, x_t, wl_{t+1})$ that govern how the intended job location d_t^w translates into the realized one wl_{t+1}. Let $v^w(wl_t, rl_t, x_t, d_t^w)$ denote the expected choice-specific value corresponding to the particular choice of job location d_t^w. We have

$$v_t^w(wl_t, rl_t, x_t, d_t^w) = \sum_{wl} \pi_t(d_t^w, wl_t, x_t, wl) EV_t^r(wl_t, rl_t, wl, x_t). \tag{7}$$

Now recall that $\epsilon_t^w \in \mathbb{R}^{R+1}$ are the stochastic components corresponding to the choice of work location, with additional voluntary choice of non-employment. Similar to the residential location choice, let $EV_t^w(wl_t, rl_t, x_t)$ be the ex ante expected value for an individual who has not learned the work location shocks $\{\epsilon_t^w(d_t^w)\}$ yet. Under the assumption that the shocks have an extreme value distribution with scale parameter σ_w, $EV_t^w(wl_t, rl_t, x_t)$ is given by the log-sum formula

$$\begin{aligned} EV_t^w(wl_t,&rl_t, x_t) \\ =&\sigma_w \log\left(\sum_{d^w} \exp\left\{\sum_{wl} \pi_t(d^w, wl_t, x_t, wl) EV_t^r(wl_t, rl_t, wl, x_t)/\sigma_w\right\}\right). \end{aligned} \tag{8}$$

Similarly, we have the usual multinomial logit choice probability for the choice of work location

$$P_t^w(d_t^w|wl_t, rl_t, x_t) = \frac{\exp\{v_t^w(wl_t, rl_t, x_t, d_t^w)/\sigma_w\}}{\sum_{d^w} \exp\{v_t^w(wl_t, rl_t, x_t, d^w)/\sigma_w\}}. \tag{9}$$

After accounting for the transition probabilities $\pi^x(x_t, x_{t+1})$ of the non-location state variables, which we assume are independent of both the stochastic shocks $\epsilon_t = (\epsilon_t^w, \epsilon_t^r)$ and the labor market probabilities $\pi_t^n(d_t^w, wl_t, x_t)$ and $\pi_t^k(wl_t, x_t)$, we have by the definition of the expected value function

$$EV_t(wl_{t+1}, rl_{t+1}, x_t) = \sum_{x_{t+1}} \pi^x(x_t, x_{t+1}) EV_{t+1}^w(wl_{t+1}, rl_{t+1}, x_{t+1}). \tag{10}$$

Combining equations (5), (7) and (10), we obtain a Bellman operator in expected value functions that maps $EV_{t+1}(wl_{t+2}, rl_{t+2}, x_{t+1})$, that enters in the shifted one period forward equation (5), into $EV_t(wl_{t+1}, rl_{t+1}, x_t)$[8].

The computational algorithm for solving the model is straightforward. Because the model is formulated in finite horizon, this reduces to a backward induction calculation starting at the maximum possible age T. For each period t we compute the expected value functions $EV_t(wl_{t+1}, rl_{t+1}, x_t)$, and the corresponding choice probabilities $P_t^w(d_t^w|wl_t, rl_t, x_t)$, and $P_t^r(d_t^r|wl_t, rl_t, wl_{t+1}, x_t)$ that serve as the basis for formulating the likelihood function.

Suppose we have already computed the expected value function $EV_{t+1}(wl_{t+2}, rl_{t+2}, x_{t+1})$ for all possible values of the states $(wl_{t+2}, rl_{t+2}, x_{t+1})$ at age $t+1$. On iteration t we loop over all possible end-of-period combinations of locations (wl_{t+1}, rl_{t+1}) and over all non-location

[8]Writing down the complete Bellman operator is straightforward, but we do not do that here for space considerations.

Table 1: Non-location state variables including household types that enter x_t.

Symbol	Description	Possible Values
cs_t	Number of children at home	0 no children 1 1 or more children
ms_t	Marital status	0 single 1 married/cohabitating
edu_t	Education (school) type	0 Less than medium cycle education 1 Medium cycle education (BA) 2 Long cycle education (master/PhD)

Notes: The table lists the non-location state variables entering x_t. In addition the value function depends on the beginning of the period work and residence locations (wl_t, rl_t), and the expected value function depends on the realized end-of-period work and residence locations (wl_{t+1}, rl_{t+1}).

states x_{t+1}. In each such point we then use equation (5) to compute the inclusive values of the different work locations wl_{t+2}, $EV^r_{t+1}(wl_{t+1}, rl_{t+1}, wl_{t+2}, x_{t+1})$, and the probabilities of location choices $P^r_t(d^r_t|wl_t, rl_t, wl_{t+1}, x_t)$ given by equation (6). Then we compute the d^w_{t+1} choice-specific values $v^w_{t+1}(wl_{t+1}, rl_{t+1}, x_{t+1}, d^w_{t+1})$ using equation (7), and the accompanying work location choice probabilities $P^w_t(d^w_t|wl_t, rl_t, x_t)$ using equation (9). After that, using the equations (8) and (10), we compute the period t expected value function $EV_t(wl_{t+1}, rl_{t+1}, x_t)$. Once $EV_t(wl_{t+1}, rl_{t+1}, x_t)$ is computed for all states $(wl_{t+1}, rl_{t+1}, x_t)$, the period t iteration is complete, and the algorithm moves to period $t-1$.

4.4 State space dynamics

There is always an awkwardness about formulating a discrete time model with actual data where transitions occur in continuous time. The discrete time model assumes decisions are made at specific instants in time: i.e. at the start of each period where the period in our setup be one year. We will defer a discussion of how to best match the actual data to the model when the precise state change date is not clear. But for our discussion suppose we have data on the state of an individual at the start of each year, i.e. on January 1st.

Table 1 lists all the state variables in the model that we include to control for the heterogeneity among the households. These variables enter the non-spatial part of the state vector x_t, and together with the two location variables form the full state vector.

As mentioned above, the transitions of the non-spatial state variables are governed by the transition probability $\pi^x(x_t, x_{t+1})$ which we describe in details below. However, a part of the state space is non-time-varying and constitute the *types* of households. We assume a finite number of types, and note that because these do not change during the backwards induction, the solution algorithm can in principle solve separate dynamic programming problems for separate types in parallel, thus assuming availability of the appropriate number of computing cores, without increasing the overall computational load. The time-invariant *type* variable is education (schooling) type edu_t and marital status ms_t[9], while children status cs_t evolve as an independent first order Markov processes with transition probabilities defined below. The potenrtial outcomes of these non-spatial (exogenous) state variables are listed in Table 1

The evolution of children status depends on age. A trick to reduce computation is to assume that the number of children can maximally change by one every year. Obviously, this fails in case of twin births and couple formation where the spouse has more than one child, but in the data we do not distinguish between having one or more children. In sum, cs_t follows

$$cs_{it+1} \sim \mu_{cs}(\cdot|cs_t, age_t). \tag{11}$$

The transitions of children are estimated separately in a first step, and the transition probability of the exogenous part of the state space vector $\pi^x(x_t, x_{t+1})$ is given by (11).

4.5 Specification of the utility function

In Section 4.3 we used the general form of the current period utility $u(wl_t, rl_t, wl_{t+1}, rl_{t+1}, x_t)$, only specifying its dependence on both current and realized work and residence locations, and the non-spatial variables. In this and next sections we give a complete specification of the utility function and job probabilities, starting with the "direct utility" specification which depends on the size of the house that the household occupies. To arrive on the final specification of $u(wl_t, rl_t, wl_{t+1}, rl_{t+1}, x_t)$ which is independent of the house size, in the next

[9]This could could be extended by a permanent income type $perminc_t$, and the propensity to move type uh_t to reflect the unobserved heterogeneity in the population. This is left for future work.

section we express the demand for housing using the first order condition of the static choice of house size, which is then plugged back into the utility function.

To help the exposition, we first describe the parts of the utility function, and specify them fully one by one afterwards. The utility of any location choice can generally be written as the sum of the following components (suppressing arguments and indices)

$$u = u_m - u_w + u_h + \underbrace{amenities - swcost_r^p - ttimecost}_{u_o}, \tag{12}$$

where u_m is the *monetary utility* (income net of housing expenditures), u_w is *disutility of work* which is equal to zero when $wl_{t+1} = \emptyset$, u_h is the *housing utility* obtained from the utilization of a chosen home size, *amenities* reflects the *regional-specific* attractiveness of housing options, $swcost_r^p$ is the *psychological costs* of changing the location of residence, and *ttimecost* is the *cost of commuting* between the chosen locations of work and residence. According to our timing convention (described in Section 4.1), all the house and regional characteristics correspond to the chosen location rl_{t+1}, because it is the location enjoyed during period t, after the instantaneous moving phase in the beginning of the period has taken place.

First, consider the u_m component. It can be expressed as a product of the marginal utility of money $\kappa(inc_t)$ (which depends on household income), and the consumable earnings, which are given by the difference between the household income and the cost of maintaining the house ($hcost_t$). We have

$$u_m = \kappa(inc_t)\big(inc_t - hcost_t\big), \tag{13}$$

where inc_t denotes household income in period t.

We assume the following functional form for the marginal utility of money

$$\kappa(inc_t) = \kappa_0 + \kappa_y inc_t. \tag{14}$$

Assuming that $\kappa_y < 0$, we have linearly decreasing marginal utility of money, implying that richer households will be less sensitive to housing prices and moving costs. In the absence of a wealth state variable and a consumption/savings choice in the model, marginal utility subsumes all effects of the credit constraint or availability of mortgage.

Household income $inc_t = inc_t(wl_t, wl_{t+1}, x_t)$ is modeled by a set of Mincer-type equations that include age as a personal characteristic, and are estimated separately for all regions and education groups to reflect the regional and skill-specific variation. In addition, we introduce a wage penalty to being non-employed in previous period ($wl_t = \o$), and we entitle the currently unemployed ($wl_{t+1} = \o$) with unemployment benefits or pension income. Conditional on being employed ($wl_{t+1} \neq \o$), household income for individual i in period t is modelled as

$$\log(inc_{it}) = \delta_0 + \delta_{age} age_{it} + \delta_{age^2} age_{it}^2 + \delta_u \mathbb{1}_{\{wl_{it}=\o\}} + \xi_{it}, \tag{15}$$

where ξ_{it} is the idiosyncratic error component and all parameters vary by work region and education group. For non-employed persons we implement the following specification of non-employment income on age dummies:

$$\log(inc_{it}) = \sum_{t=t_0}^{T} b_t + \nu_{it}, \tag{16}$$

for each education group, where ν_{it} is a random error term.

The probabilities of getting a new job or keeping the existing job were introduced in Section 4.2. Below we present the functional forms, starting with the probability of getting a new job which is defined as

$$\begin{aligned} \pi_t^n(d_t^w, wl_t, x_t) = \Bigg[1 + exp\Bigg(- \Big(\beta_0^{\pi(new)} + \beta_a^{\pi(new)} age_t + \beta_{unemp}^{\pi(new)} \mathbb{1}_{wl_t=\o} \\ + \beta_{jobdensity}^{\pi(new)} jobdensity(d_t^w) \\ + \sum_{k=1}^{2} \big(\beta_s^{\pi(new)}(k) \mathbb{1}_{edu_t=k}\big)\Big)\Bigg)\Bigg]^{-1}, \end{aligned} \tag{17}$$

where $jobdensity(d_t^w)$ is an index of type edu_t jobs in region d_t^w. By allowing the probability of landing a new job to depend on job density, then when $\beta_{jobdensity}^{\pi(new)} > 0$ the individual is more likely to receive a job offer from a region with more jobs. This helps the model predict how attractive each work region is. Admittedly, there is something awkward about this specification as long as we do not model the equilibrium on the labor market, since we attribute a high job density to a high fixed supply of jobs while in reality it is an interplay between supply and demand. We assume the individuals only cares about and searches for jobs of their own skill type to capture the heterogeneity in job moving behaviour. The probability of keeping one's current job is defined by

$$\pi_t^k(wl_t, x_t) = \left[1 + exp\Bigg(- \Big(\beta_0^{\pi(keep)} + \beta_a^{\pi(keep)} age_t + \sum_{k=1}^{2} \big(\beta_s^{\pi(keep)}(k) \mathbb{1}_{edu_t=k}\big)\Big)\Bigg)\right]^{-1}. \tag{18}$$

Finally, we allow for disutility of work u_w through the fixed constant, c_{work}, which is relevant when $wl_{t+1} \neq \emptyset$.

Amenities of home regions are modelled as region-specific constants. Another approach would be to model amenities as a function of region-specific observables such as crime rates, nature, restaurants etc., but given that most of these variables are regional-specific and time constant we use a fixed effects approach:

$$amenities(rl_{t+1}) = \sum_{rl=1}^{R} \boldsymbol{\alpha}_{rl} \mathbb{1}_{\{rl_{t+1}=rl\}}, \tag{19}$$

where α^{rl} is a vector of coefficients for each region.

The psychological moving cost $swcost_r^p$ is a function of the family characteristics, age and education. We use the following specification

$$swcost_r^p(x_t) = \mathbb{1}_{\{rl_t \neq rl_{t+1}\}} \big[\gamma_0 + \gamma_{age} age_t + \gamma_{ms} ms_t + \gamma_{cs} cs_t + \sum_{k=1}^{2} \phi_{s,k} \mathbb{1}_{\{edu_t=k\}}\big], \tag{20}$$

which reflects the fact that the propensity to move changes with the family situation and is different at different stages of life.

The costs of commuting between rl_{t+1} and wl_{t+1} are assumed to be proportional to the exogenous travel time between the work and home locations. Hence, we have

$$ttimecost = \eta_{ttime} ttime(rl_{t+1}, wl_{t+1}) \tag{21}$$

where the function $ttime(rl_{t+1}, wl_{t+1})$ denotes the travel time between work location wl_{t+1} and residence location rl_{t+1}.

When specifying the demand and utility of housing, we note that regional-specific price of housing is approximately linear in home size measured in square meters of floor space. It is therefore natural to specify housing demand (size of home) $h(rl_{t+1}, x_{it}; P^h(rl_{t+1}))$ in residential region rl_{t+1} as a function of the regional-specific housing price $P^h(rl_{t+1})$ (expressed as an equivalent annual rental price) per square meter. Housing costs $hcost_t$ are thus assumed to be proportional to the equilibrium per square meter prices $P(rl_{t+1})$ through the parameter ψ_{uc}. This translates housing prices into an annual user cost, and also reflects mortgage expenses and housing taxes. Hence housing costs are given by

$$hcost_t(rl_{t+1}, h_{t+1}) = \psi_{uc} P(rl_{t+1}) h_{t+1}, \tag{22}$$

where according to our timing convention rl_{t+1} denotes the house occupied during period t.

The demand for housing also depends on individual characteristics such as household size and income. This for example reflects that richer people can buy relatively more square meters and others less, and that larger families may substitute space for location. We define the utility u_h of living in a house as a quadratic polynomial of its size h_t with heterogeneous coefficients

$$u_h = \Phi(x_t) h_{t+1} + \frac{1}{2} \phi_{h2} \, h_{t+1}^2, \tag{23}$$

where $\phi_{h2} < 0$ governs the degree of diminishing returns to house size and $\Phi(x_t)$ is a heterogeneous parameter, which affects the baseline marginal utility of housing

$$\Phi(x_t) = \phi_0 + \phi_{age} age_t + \phi_{ms} ms_t + \phi_{cs} cs_t. \tag{24}$$

Given the form of the utility function specified by equations (12)-(18), the part of utility that is dependent on home size is equal to

$$\tilde{u}_h = \Phi(x_t)h_{t+1} + \frac{1}{2}\phi_{h2}\, h_{t+1}^2 - \kappa(inc_t)hcost_t(rl_{t+1}, h_{t+1}). \tag{25}$$

From the first order conditions for (25) and the specification of housing cost in (22) it follows that the optimal choice of the house size[10] is given by

$$h_{t+1} = \frac{\kappa(inc_t)P(rl_{t+1})\psi_{uc} - \Phi(x_t)}{\phi_{h2}}. \tag{26}$$

Substituting expression (26) back into the utility function defined in equations (12)-(25), we obtain the final specification of the indirect utility function $u(wl_t, rl_t, wl_{t+1}, rl_{t+1}, x_t)$.

5 Structural Estimation

This section describes the estimation strategy applied to the theoretical model. We estimate the model sequentially in three steps: i) we estimate the parameters governing the wage equations and transition probabilities of the children state; ii) we estimate a reduced form housing demand equation, and iii) we estimate the remaining structural parameters by maximum likelihood applying the parameters obtained in i) and ii). Below we go through each step in detail and discuss identification of parameters.

First a note on our data sampling. In Section 3 we provide descriptive statistics using full population register data for the period 1998-2010, yet for the estimation we focus on years 2005-2010. We do this to work with a relatively homogeneous subsample while also retaining relevance in terms of policy guidance. Further, even though we do observe each individual's choice $d_{i,t} \equiv \{rl_{i,t+1}, wl_{i,t+1}, h_{i,t+1}\}$ and state $s_{i,t}$ on an annual basis, we pool the data over all the years in estimation in the current version of the estimation.

For time varying variables such as house prices, we average them over the same period. Hence, we assume that all choices made by households are made considering only the average

[10] At the time immediately after all moving has finished since, as is clear from the Section 6, the equilibrium of the model is Markov perfect and higher demand in a region drives up prices.

housing prices and amenities over this period. Since local amenities are fixed throughout the estimation period, we cannot separately identify the effects of observed regional amenities and regional fixed effects. A more subtle note here is that since both local house prices and local amenities are regional-specific in the pooled sample, joint identification of the utility coefficients of the time-invariant local amenities and marginal utility of money can only be accomplished through individual-level taste variation for housing and individual-level variation in income. However, since both income and the demand for housing varies within each region we are able to simultaneously pin down amenities and marginal utility of money along with the remaining parameters of the model.

5.1 Wage equations and transition probabilities

The estimation of transition probabilities for children status, μ_{cs}, is performed non-parametrically on the pooled data as the share of individuals within each age-children cell who is observed in each possible transition. Since in the current implementation, cs only takes values 0 and 1 (at least one child living at home), this comes down to four possible transitions at each age.

In order to capture regional differences in both wage level and its age gradient, we estimate the coefficients of the wage offer equation in (15) separately for each combination of region and education level. We use the equivalent full-time income for each individual and condition on observed employment. The estimates are presented in Appendix C.

Similarly, we estimate the education level-specific equation for non-employment income in (16). Until retirement age, non-employment income will mainly consist of unemployment benefits. After the usual retirement age, income sources are more mixed as one will receive both public pension, private pension savings and possibly labor income from part time employment.[11]

[11] We do not allow for regional differences in non-employment income, even though these are indeed observed due to differences in savings. Yet, since an individual would not be able to change her savings by moving, we abstract away from differences in average regional savings. The result is of course that the returns to income received while working is downward biased in rich areas and upward biased in poor areas. Estimates of the non-employment income regressions are not presented but available upon request.

5.2 Housing demand

As mentioned in Section 4.1, we assume that within each region and in each time period households can freely adjust the size/quality of their home. This is equivalent to having no cost of moving within the region to the house of optimal size. Moreover, we abstract away from any savings, including in home equity, and let households consider only the "square meter rental costs" that pertains to homes in each region through local prices. Both of these assumptions allow for the optimal amount of housing to be separable from the dynamic choice of location and expressed as the solution to a static subproblem that enters into the indirect instantaneous utility described in Section 4.5. This greatly reduces the computational burden, effectively allowing the structural estimation of the model to be carried out.

The fact that the solution to the static housing size problem specified in (26) is detached from the dynamic location choice allows us to estimate a scaled version of it in a separate step before turning to the dynamic model. Hence, using the pooled micro data we estimate the following demand equation for housing by OLS

$$\begin{aligned} h_{it+1} = {} & \tilde{\phi}_0 + \tilde{\phi}_{age} age_{it} + \tilde{\phi}_{ms} ms_{it} + \tilde{\phi}_{cs} cs_{it} \\ & - \tilde{\kappa}_0 [\psi_{uc} P(rl_{it+1})] + \tilde{\kappa}_y [inc_{it} \times \psi_{uc} P(rl_{it+1})] + \varrho_{it}, \end{aligned} \tag{27}$$

where ϱ_{it} is a random error (see also equation (26)).

Note that the parameters $\tilde{\phi}$ and $\tilde{\kappa}$ in the reduced form demand equation in (27) are proportional to the structural parameters that index marginal utility of money $\kappa_{(\cdot)}$ and heterogeneous housing utility parameters in $\Phi_{(\cdot)}$, but scaled by $-1/\phi_{h2} > 0$ and $-\psi_{uc}/\phi_{h2} > 0$ respectively. We identify ϕ_{h2} and ψ_{uc} in conjunction with the remaining structural parameters using the cross-equation restrictions implied by the housing demand equation and the location choice model. When estimating these structural parameters, the reduced form estimates of parameters are then kept fixed during the structural estimation of the location choice model, and only *rescaled* using the values of the structural parameters ψ_{uc} and ϕ_{h2}. This two-step procedure significantly reduces the dimensionality of the maximum likelihood problem when estimating the full model.

5.3 Maximum likelihood estimation of structural parameters

Having obtained estimates for children transitions, wage equations and scaled housing demand we estimate the remaining structural parameters, θ, by maximum likelihood. To recount, θ includes parameters indexing probability of getting a new job, (17), probability of keeping current job, (18), marginal utility of money, (14), housing costs, (22), utility values of the amenities, (19), psychological costs of moving residence, (20), travel time costs, (21), the disutility of work, c_{work}, and the degree of diminishing returns to house size, ϕ_{h_2}. We fix the discount factor to $\beta = 0.95$.

The likelihood function is derived from the choice probabilities for work and home location decisions given in (6) and (9). Because we assume perfect control for residential location, the latter can be directly evaluated at the data, giving the likelihood of the observed location of residence. To calculate the likelihood of the observed work location, however, we have to integrate out the likelihood over the possible choices and only write the likelihood in terms of observed *work location transitions*, i.e. as transition probabilities from state wl_t to wl_{t+1}.

Observing a transition wl_t to $wl_{t+1} = wl_t$ could have resulted from both an individual deciding to keep their job, and being successful (with probability $\pi_t^k(wl_t, x_t)$), and an individual trying to find a new job d_t^w and being unsuccessful (with probability $\big(1 - \pi_t^n(d_t^w, wl_t, x_t)\big)\pi_t^k(wl_t, x_t)$). Observing a transition wl_t to $wl_{t+1} \neq wl_t$ could have resulted only from an individual deciding to move jobs and being successful (with probability $\pi_t^n(wl_{t+1}, wl_t, x_t)$).

The above two cases also apply for $wl_t = \emptyset$, but $wl_{t+1} = \emptyset$ and $wl_t \neq \emptyset$, i.e. transition to unemployment, may happen in three different scenarios. First, with probability $\big(1 - \pi_t^n(d_t^w, wl_t, x_t)\big)\big(1 - \pi_t^k(wl_t, x_t)\big)$ an individual could have unsuccessfully tried to transition to a job d_t^w, and at the same time has been displaced. Or, with probability $1 - \pi_t^k(wl_t, x_t)$ an individual could have tried to keep job wl_t, yet being unsuccessful and displaced. Or finally, an individual could have voluntarily chosen to quit working, $d_t^w = \emptyset$.

Recall that $\pi_t(d_t^w, wl_t, x_t, wl_{t+1})$ summarizes the work transition probabilities as a function of the intended work location. The contribution to the likelihood for an individual who is in *observed* work location wl_t and residential location rl_t at time t and in *observed* work location wl_{t+1} and residential location rl_{t+1} at time $t+1$ is

$$L_t(wl_t, rl_t, wl_{t+1}, rl_{t+1}, x_t) =$$
$$P_t^r(rl_{t+1}|wl_t, rl_t, wl_{t+1}, x_t) \cdot \sum_{d^w} P_t^w(d^w|wl_t, rl_t, x_t)\pi_t(d^w, wl_t, x_t, wl_{t+1}). \tag{28}$$

The full log-likelihood is constructed from individual likelihoods in the standard way by collecting the individual likelihood contributions and the objective of the maximum likelihood estimation is thus

$$\text{argmax}_\theta \frac{1}{N}\sum_i \sum_t \{ \log P_t^r(rl_{it+1}|wl_{it}, rl_{it}, wl_{it+1}, x_{it}; \theta)+$$
$$\log \sum_{d^w} P_t^w(d^w|wl_{it}, rl_{it}, x_{it}; \theta)\pi_t(d^w, wl_{it}, x_{it}, wl_{it+1}; \theta)\}, \tag{29}$$

where N is the number of individuals. To estimate the structural parameters we proceed in the spirit of the Nested Fixed Point (NFXP) algorithm by Rust (1987, 1988) and solve the model via backwards induction for each evaluation of the likelihood function.

5.4 Identification

We now consider how the parameters of the model are identified from data. Starting with the first-step estimation of housing demand, recall that Section 3.3.3 clearly demonstrates variation in housing demand by age, number of children and marital status. As this variation is present within regions as well we have identification of the parameters in $\tilde{\Phi}$ which were scaled by $-\phi_{h2}$. The same goes for the parameters governing marginal utility of money in $\tilde{\kappa}$. The baseline marginal utility of money κ_0 is identified from the spatial variation in house prices together with individual-level variation in housing demand and income. Further, the income dependence in marginal utility of money, κ_y is identified from the clear sorting of high-income households into larger homes and more expensive regions.

It follows from (26) and (27) that we can only identify the parameter for diminishing utility of housing, ϕ_{h2}, within the dynamic model of location decision. The dynamic location choice involves a trade-off between home size, value of amenities and commuting time so the fact that we observe households substituting between locating in regions of high amenity levels in return for smaller homes than in low amenity regions allows us to determine ϕ_{h2}.

The user cost of housing, ψ_{uc}, and marginal utility of money, κ, can be separately identified using the variation in location decisions, house prices and wages across regions. For a given marginal utility of income, the sensitivity of individuals' location decisions to the spatial variation in house prices provides identification of ψ_{uc}. In that sense ψ_{uc} can also be thought of as a factor that distinguishes individuals' marginal utility of wage income from marginal (dis)utility of house prices.

The disutility of work, c_{work}, is identified through both the variation in labor market participation across education groups and through the participation over the life cycle. High skill workers have both higher wages and higher participation rates, implying that the opportunity cost of not working must be positive, thus $c_{work} > 0$. The same effect occurs within education groups as the wage offer declines after a certain age which coincides with increasing propensity to not work (i.e. retire).

The parameters that index moving costs, $\{\gamma_0, \gamma_a, \gamma_c, \gamma_{ms}, \gamma_s\}$, are easily identified since the propensity to relocate home differs substantially along the age, children, marital and schooling dimensions. The variation in moving propensity along age, kids and education is evident in the graphs of Section 3.3 and corresponding graphs are shown in the results in Section 7 for education groups. Higher age, presence of kids and lower education all reduces the likelihood of moving.

The coefficients with local amenities in $\alpha_{rl}(\cdot)$ are identified by observing that at a given level of income and commute distance, households are willing to pay a higher square meter price in one region compared to another. The only justification for such behavior is a higher amenity level.

The uncertainty of the job search process implies that large work regions in terms of number of jobs are more attractive than small regions conditional on the wage offer. The parameter $\beta^{\pi(new)}_{jobdensity}$ is therefore determined by differences in transitions of a given household type into regions that offer similar wages but have different sizes. Since we also observe that there is a negative gradient in job relocations over the life cycle without a corresponding drop in wage differences, we get identification of the age component of job search. Similarly, less schooling and past unemployment will affect job transitions negatively as long as the relative differences in income across regions do not decrease one for one with these measures. We do not try to identify any local effects in the probability of keeping a job. Hence, observed transitions into non-employment across regions that offer similar wages for a given household type (in the age and education dimension) yield identification of keep probabilities.

Conditional on a choice of work location, wl, we observe a decreasing probability of choosing a home location rl as the distance to wl increases. This relationship pins down the cost of traveltime η_{ttime}.

6 Solving for Equilibrium House Prices

We take a short run perspective and assume a fixed supply of housing and thus abstract from the longer run dynamics where new houses are built in response to changes in house prices. We also abstract from equilibrium formation in the labor market, and ignore that firms in reality may change labor demand in their locations (and thus the number of jobs offered in different locations) in response to changes in local labor supply. Hence the equilibrium object we solve for is housing/rental prices, whereas wages, job arrival and dismissal rates are taken as given and housing supply is assumed fixed in the short run.

In equilibrium we assume that prices have adjusted so that the total demand for housing measured in square meters equals the supply in each residential region. Thus, when solving for the housing market equilibrium, the R-dimensional vector of regional square meter prices $P^h = (P^h(1), .., P^h(R))$ is set to equate the inelastic, exogenously fixed supply $S_t(rl)$ of total square meters of housing to the demand for the available square meters $D_t(rl, P^h)$ in each residential region $rl = \{1, \ldots R\}$. For the supply, we simply aggregate the individual-level demand for observed square meters of housing h_{it} for people who already live in region

$rl_{it} = rl$ at the beginning of each period t

$$S_t(rl) = \sum_{i=1}^{N} h_{it} \mathbb{1}(rl_{it} = rl) \tag{30}$$

where $\mathbb{1}$ is the indicator function.

The regional demand for housing $D_t(rl, P^h)$ is calculated as the *expected demand* by taking a population average of housing demand weighted by choice probabilities of either staying or moving to region rl at the end of period t. To obtain demand, we start by simulating N individual states by drawing from observed states in the dataset with replacement. We then simulate a work location *outcome*, wl_{t+1}, using the decision rule P_t^w and job transition probabilities π_t such that we can condition on these in the computation of demand below:

$$D_t(rl, P^h) = \sum_{i=1}^{N} h(rl, x_{it}; P^h(rl)) \Pi_t(rl|wl_{it+1}, rl_{it}, x_{it}; P^h), \tag{31}$$

where $\Pi_t(rl|wl_{it+1}, rl_{it}, x_{it}; P^h)$ is the probability that an individual in state $s_{it} = (wl_{it+1}, rl_{it}, x_{it})$ chooses to live in region rl given the vector of regional house prices, P^h and simulated work location wl_{it+1}. Π_t is given by the right hand side of (6), but here we have added P^h as an argument to signify its dependence on house prices.

The resulting simulator for demand is in principle not smooth given that we have simulated a work location *outcome*, wl_{t+1} using a simple accept/reject simulator. However, since the conditional demand for residence, $\Pi_t(rl|wl_{it+1}, rl_{it}, x_{it}; P^h)$, is smooth in the vector of housing prices and employment probabilities, we still found it smooth enough to use gradient-based methods to calculate equilibrium. We calculate the house price equilibrium by arraying all the excess demand equations to have a system of R excess demand equations (for the housing market) in R unknowns and solve for the R-dimensional price vector P^h using Newton's method.

The short run equilibrium concept is imposed for simplicity. To work with a long run equilibrium notion that endogenizes the supply of housing, we would need data on zoning regulations and decisions by home builders and developers where to build more in different

Table 2: First Stage Parameter Estimates, Housing Demand

Variable (parameters)	Coeff. Estimates	Standard Error	t-statistic
Const ($-\phi_0/\phi_{h2}$)	122.3154	0.05752	2126.3
Married ($-\phi_{ms}/\phi_{h2}$)	19.4172	0.01517	1279.7
Children ($-\phi_c/\phi_{h2}$)	13.6033	0.01615	842.2
Age ($-\phi_a/\phi_{h2}$)	0.5824	0.00059	983.6
Price pr. sqm ($\kappa_0\psi_{uc}/\phi_{h2}$)	-304.1712	0.21142	-1438.7
Price pr. sqm $\times$ income ($\kappa_y\psi_{uc}/\phi_{h2}$)	21.3753	0.02827	756.1

regions. Finally, commuting times/costs are potentially something to endogenize too, including in the short run. If the counterfactual equilibrium results in changed location patters, the resulting utilization of the road network will change as well and thereby affect congestion and commuting times. Future work will focus on these more involved specifications.

7 Results

We start by presenting model fits and parameter estimates from the first-stage income and housing demand equations and then move to the remaining parameter estimates from the structural location choice model. Using the estimated model to solve for equilibrium prices, we analyze the in-sample fit of computed equilibrium house prices compared to observed house prices. Finally, we conduct counterfactual policy experiments where we increase housing supply and commute time and compare the predicted responses in terms of residential sorting and job location.

7.1 Parameter estimates and model fit

The parameters estimates are provided in Tables 2 through 6[12]. Table 2 presents the estimates obtained from the housing demand regression in (27). Note that both annual income and housing price per square meter are measured in units of 100,000 DKK and we therefore use this unit in the following examples. A slight complication is that annual income is recorded before taxes while housing expenses obviously must be paid after taxes. Therefore, the implicit willingness to pay for housing, amenities and commuting will be measured in pre-tax

[12]Parameter estimates from income regressions for employment regions available in Appendix C. Parameter estimates for non-employment not shown but available upon request.

Table 3: Curvature Parameter of Housing Demand and User Cost

	Coeff. Estimates	Standard Error	t-statistic
Coef. on h^2, ϕ_{h2}	-0.0007	0.00000	-865.0
User cost housing, ψ_{uc}	0.2466	0.00139	177.9
κ_0	0.863		
κ_y	-0.061		

income rather than actual disposable income. To avoid this issue, one would have to model the tax system on top of the wage equations, which we have deferred from in current version.

The coefficients of the reducd form housing demand presented in Table 2 have reasonable magnitudes and expected signs. Recall from Section 5 that our estimation strategy only allowed for identification of scaled parameters in the first step housing demand. We can however deduce that demand is increasing as a function of age and household size, and couples live in homes that, on average, measure 19.4 m^2 more than singles. Having children living at home is associated with a 13.6 m^2 larger dwelling. Housing demand is decreasing in prices as $\frac{\kappa_0 \psi_{uc}}{\phi_{h2}} < 0$, yet to a lesser extent for richer individuals since the term interacted with income is positive.

During the sample period, the average price per square meter was $26,661$ DKK in Copenhagen. Hence, individuals choosing to live in the Copenhagen municipality will on average demand $21.375 * 0.266 = 5.7$ more square meters of housing for each additional 100,000 DKK of individual annual income. Similarly, an individual with an income of 500,000 DKK living in Copenhagen demands 13.7 fewer square meters of housing compared to an individual with similar income living outside the capital area (Rest of Zealand) where square meter prices are $19,704$ DKK on average, i.e. around 7,000 DKK lower than in Copenhagen municipality[13].

The parameter estimates for ϕ_{h2} and ψ_{uc} are given in Table 3. We estimate the annual user costs of housing to $\hat{\psi}_{uc} = 0.247$, ie. 24.7% of the market value. This is definitely on the high side, but there are certain factors that may explain it. First of all, as noted above there is the complication that income is recorded before-tax. Since the tax burden lies in

[13]The difference in housing demand across these two regions is computed as $(-304.17+21.37*5)*(0.26661-0.19703) = -13.73$.

Figure 9: Model fit: housing size over the life cycle

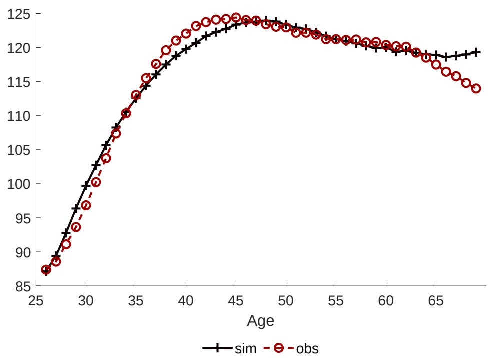

(a) Overall

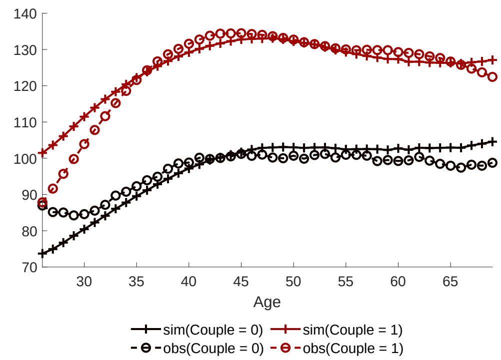

(b) By marital status

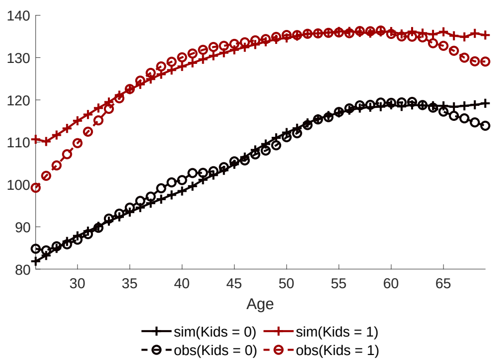

(c) By children

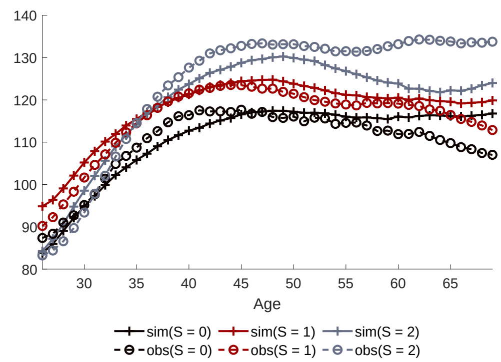

(d) By schooling

the interval 30-50%, the user cost measured in disposable income is correspondingly lower, around $12 - 17\%$. Furthermore, our estimation period 2005-2010 is mostly characterized by falling housing prices. In the standard user cost equation for housing, expected discounted capital gains reduce the user cost. If that equation truly lies in the back of people's mind when making housing purchases, then falling prices and pessimistic expectations work to increase user costs, and this might be what our estimate of ψ_{uc} is picking up.

The model fit of chosen house size over the life cycle is shown in Figure 9a. The parameters capture the change in the demand over the life cycle closely. Separating by household size both in terms of having children and having a partner the model also provides a reasonable fit. There are some challenges capturing the demand at the beginning and end of the life cycle. The same goes for demand by education groups, where there is an underprediction of demand for the highly educated and overprediction for the medium- and low-skilled at the end of the life cycle. These obstacles are likely a result of the fact that we do not allow for

Figure 10: Model fit: income and housing size by home region

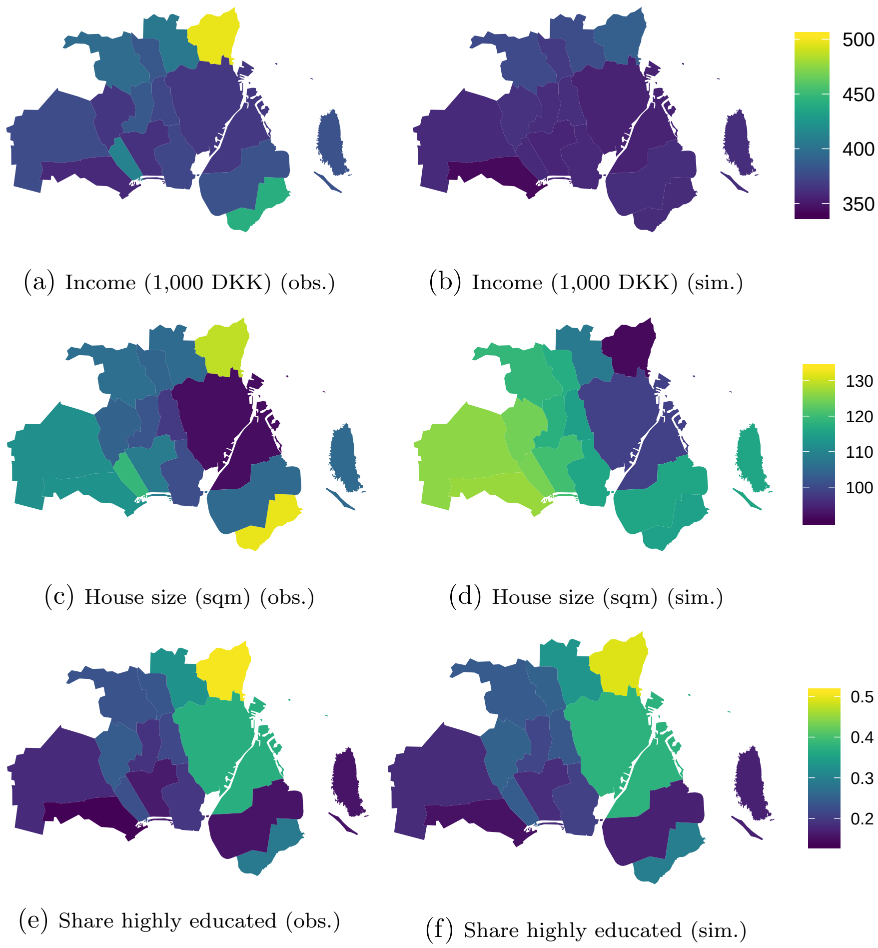

(a) Income (1,000 DKK) (obs.)

(b) Income (1,000 DKK) (sim.)

(c) House size (sqm) (obs.)

(d) House size (sqm) (sim.)

(e) Share highly educated (obs.)

(f) Share highly educated (sim.)

Note: Panels (a) and (b) show the average income in 1,000 DKK by home region. Panels (c) and (d) show the average size of homes in square meters by home region.

adjustment costs and savings to affect housing size decisions.

Using the estimates of $\phi_{h2} = -0.0007$ and $\psi_{uc} = 0.2466$ together with the reduced form estimates in housing demand given in Table 2, we can back out the parameters that index marginal utility of money. We obtain $\kappa_0 = 0.86$ and $\kappa_y = -0.061$. Despite the strong negative gradient in income, these parameters result in relatively large estimates of marginal utility of money throughout most of the income distribution. Therefore the parameters imply a strong trade-off between home size and residential location and a clear sorting by richer individuals into more attractive and expensive regions and larger houses.

Figure 10 illustrates the ability of the model to fit sorting by highly educated individuals in conjunction with the variation in average income and housing demand across regions. The distribution of highly educated is captured very well because the income equation is specifically tied to the individual's education, and income predicts the home location through marginal utility of money. For example, the model is able to predict that the share of highly educated is high in Copenhagen, Frederiksberg and Gentofte where per square meter prices are high.

Although the model captures the educational sorting quite well, it has difficulty capturing the income levels in Gentofte, Dragoer and Vallensbaek. The latter would be improved if we included more heterogeneity in the income specifications such as the lagged dependent variable and persistent unobserved heterogeneity. Without sufficient variation in income, it is also hard to explain the spatial distribution of house sizes. The house size are especially underpredicted for Gentofte and Dragoer where individuals' incomes are the highest according to Figure 10a.

The residential sorting is driven mainly by four factors: i) regional variation in house prices and regional-specific amenities, ii) individual differences in housing demand, iii) individual differences in marginal utility of money and iv) distance to local labor markets. To flexibly capture regional-specific amenities, we include fixed effects for each residential region, $\alpha_{rl_{t+1}}$. The presence of local fixed amenities help rationalize why individuals prefer to live in regions where prices are high for reasons that are not explained by factors such as better access to local labor markets. The parameter estimates are presented in Table D1 in the appendix.

Gentofte (region 5) is associated with the highest amenity level and together with Frederiksberg (region 1) these are the only regions with better amenities than Copenhagen municipality (region 0) which is the outside category. The least attractive regions in terms of amenities are Albertslund (region 9), Hoeje-Taastrup (region 11), Ishoej (region 13) and Vallensbaek (region 15). They are all located in a cluster on the south-western border of the Greater Copenhagen Area. To exemplify the magnitudes, an individual with an annual income of $500,000$ DKK would need $1.7885/(0.86 - 0.061 \cdot 5) = 3.22$, i.e. 322,252 DKK, in compensation for living a year with the amenity level of Ishoej rather than that of Copenhagen municipality.

Figure 11: Model fit: residential sorting

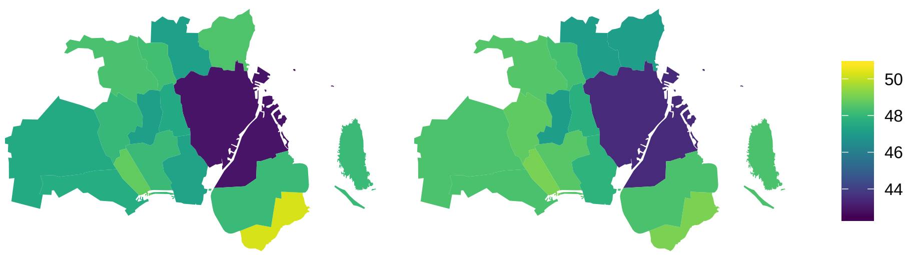

(a) Age (obs.) (b) Age (sim.)

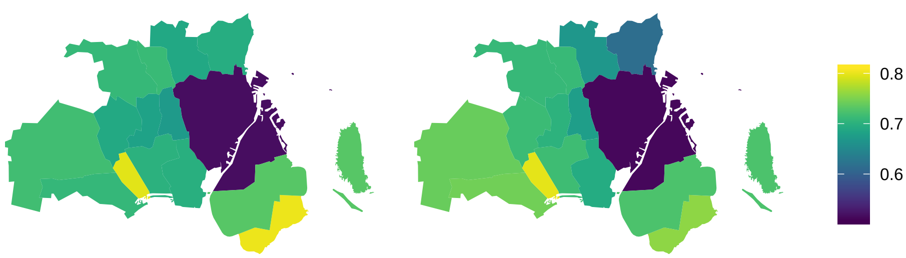

(c) Couples (obs.) (d) Couples (sim.)

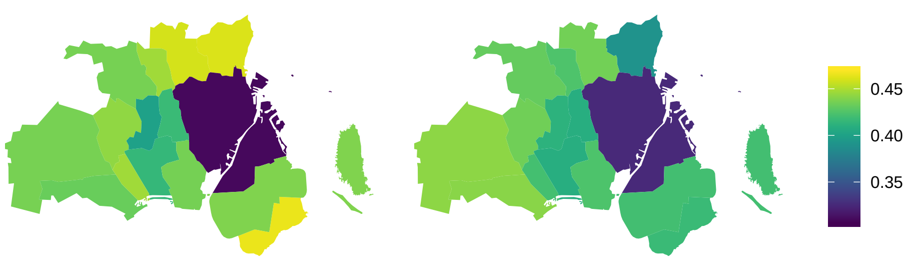

(e) Children (obs.) (f) Children (sim.)

Note: Panels (a) and (b) show average age by home region. Panels (c) and (d) show the share of couples by home region. Panels (e) and (f) show share of households with children by home region.

Figure 12: Model fit: share living in Copenhagen over the life cycle

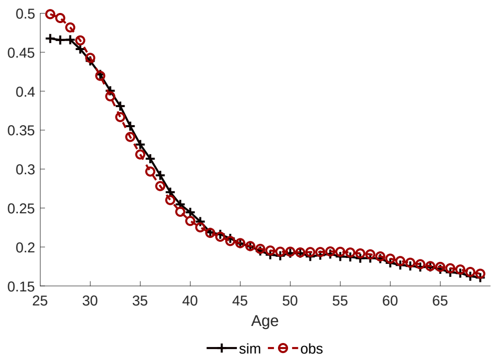

(a) Overall

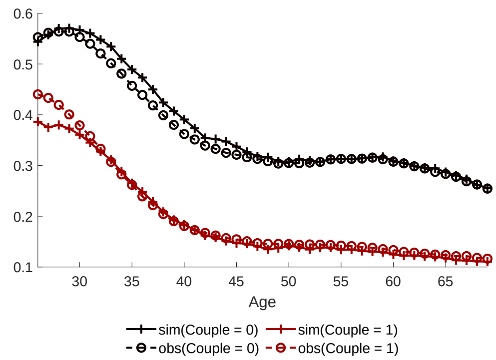

(b) By marital status

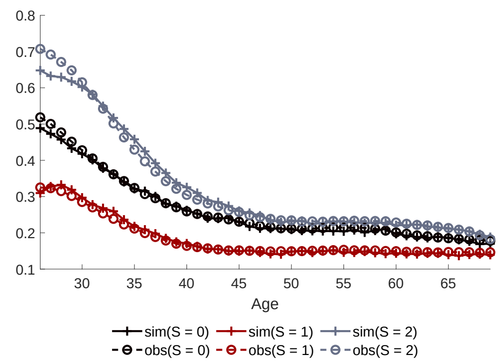

(c) By schooling

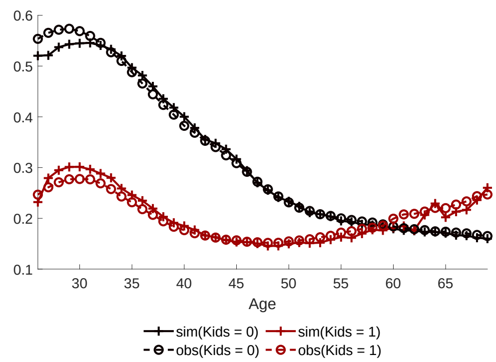

(d) By children

Figure 11 presents the model fit in terms of the residential sorting of household demographics. Starting with Figure 11a, the average age of the individual by home region is well captured. It is only slightly underpredicted in Dragoer, Gentofte and on the border between Rest of Zealand and the GCA. Looking at the share of couples in each region, the model fit also looks good. Again, Dragoer and Gentofte stand out as regions where the model underpredicts the shares the most. The distribution of families with children is less accurately captured, cf. Figure 11f. This could be improved by interacting the fixed effects by a dummy for children to pick up if the unobserved amenities of regions are valued differently by households with children.

Figure 12a shows corresponding fits for the probability of living in Copenhagen, but over the life cycle instead of the spatial allocation. The fit is very good in all respects. Only for the youngest cohorts is there a slight underprediction of the share living in Copenhagen. This is partly due to the fact that we do not model educational choice, and many higher educational institutions are located in Copenhagen. Figure 12c does indeed show that this problem is only evident for individuals with high education. It should be noted that these moments are not only driven by the estimates of amenity values, but to a large extent by the moving costs that prevent people from moving away from their initial locations.

Table 4: Utility Cost of Moving Residence

	Coeff. Estimates	Standard Error	t-statistic
Const., γ_0	1.8363	0.00921	199.4
Age, γ_a	0.0881	0.00021	420.3
Married, γ_{ms}	0.0605	0.00485	12.5
Children, γ_c	0.8212	0.00523	156.9
Schooling, γ_s (1)	0.1797	0.00553	32.5
Schooling, γ_s (2)	-0.1470	0.00545	-27.0

Table 4 displays the estimates for the parameters γ that index the utility cost associated with moving residence. Married individuals and those with children are predicted to have higher moving costs and more so as they age. Medium-skilled individuals are less likely to move, all else equal, compared to low- and high-skilled types. Individuals with highest education are more mobile.

Overall, the model fit in terms of residential moving probabilities is good according to Figure 13a. There is a slight overprediction in the start of the life cycle, especially for individuals without children and singles. The largest prediction error is found for the probability of moving to and from Copenhagen. Figure 14 shows that the general shape of the probability of moving away from Copenhagen (as a share of all individuals in our data) is captured by the model, but it underpredicts the level until the age of 45. Conditioning on those who move, Figure 14b reveals that the same problem is observed among the movers, but the magnitude is larger. As the lower panel shows, the same can be said about the share migrating to Copenhagen. A key factor left out of the model is that we ignore the obvious fact that Copenhagen is a university city. Without explicitly modeling educational choice and the dynamics of occupational career choice it is hard to explain why younger individuals with low income choose to live in Copenhagen. Other omitted factors are individual taste variation for regional-specific amenities such as bars, restaurants, child care, and school quality which can readily be included into this model at a low computational cost. Also by modelling moving

Figure 13: Model fit: share moving residential location over the life cycle

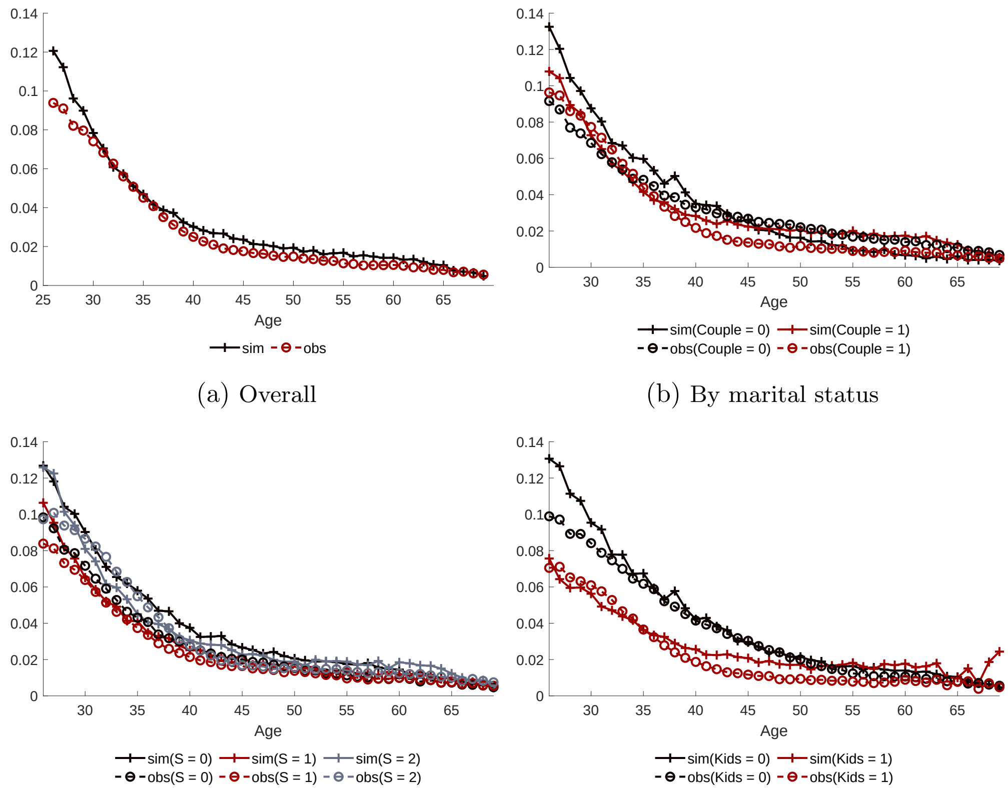

(a) Overall

(b) By marital status

(c) By schooling

(d) By children

Figure 14: Model fit: share moving residential location from and to Copenhagen over the life cycle

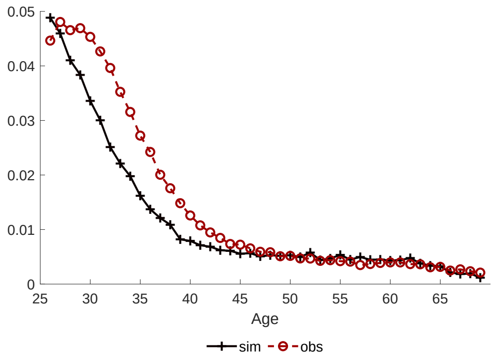

(a) From Copenhagen out of all

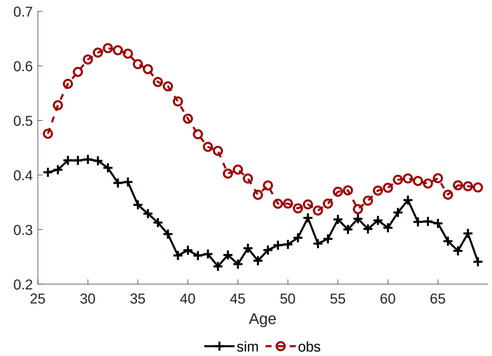

(b) From Copenhagen out of all movers

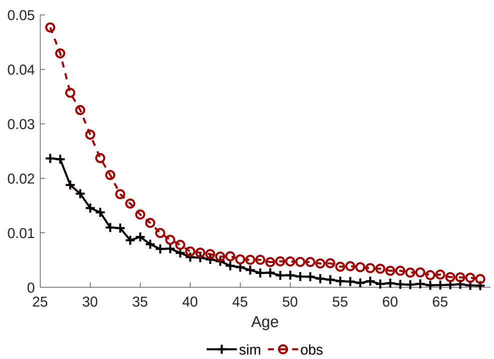

(c) To Copenhagen out of all

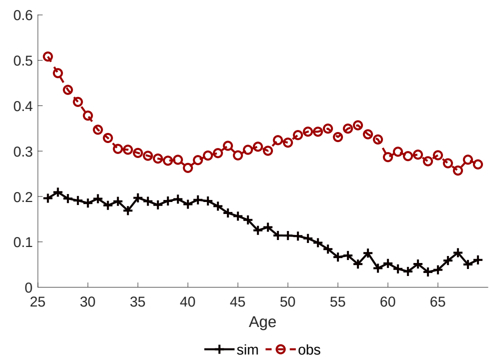

(d) To Copenhagen out of all movers

Note: Panels (a) and (c) show the share of all individuals in the data who move residential location from and to Copenhagen, respectively. Panel (b) and (d) show the share of all residential movers who move from and to Copenhagen, respectively.

costs more carefully, e.g. including unobserved heterogeneity, we may be able to predict these shares better.

We now move to the ability of the model to predict work location outcomes. Table 5 displays estimates for the parameters for the job arrival and dismissal probabilities $\pi_t^n(d_t^w, wl_t, x_t : \beta^n)$ and $\pi_t^k(wl_t, x_t; \beta^k)$ that determines the work location transition probabilities. Starting with the probabilities of keeping the job, there is a positive effect of age and higher levels of schooling. An individual who is 40 years old and has a low-level education has a $(1 + \exp(-(2.435 + 0.014 \cdot 40)))^{-1} \cdot 100 = 97.7$ percent chance of keeping the job. A similar person, who was working in $t-1$ and searches for a new job has a 55.4 percent chance of being successful and ending in Copenhagen. In Hoeje-Taastrup, where the job density for

Table 5: Job Arrival and Dismissal

	Coeff. Estimates	Standard Error	t-statistic
Probability of keeping job: $\pi_t^k(wl_t, x_t; \beta^k)$			
Const., $\beta_0^{\pi(keep)}$	2.2226	0.04122	53.9
Age, $\beta_a^{\pi(keep)}$	0.0384	0.00098	39.0
Schooling, $\beta_s^{\pi(keep)}$ (1)	0.8267	0.02178	38.0
Schooling, $\beta_s^{\pi(keep)}$ (2)	0.5677	0.01633	34.8
Probability of new job: $\pi_t^n(d_t^w, wl_t, x_t : \beta^n)$			
Const., $\beta_0^{\pi(new)}$	-0.2453	0.00617	-39.7
Age, $\beta_a^{\pi(new)}$	-0.0624	0.00014	-457.6
Schooling, $\beta_s^{\pi(new)}$ (1)	0.1455	0.00347	41.9
Schooling, $\beta_s^{\pi(new)}$ (2)	0.2580	0.00375	68.8
Job density $\beta_{jobdensity}^{\pi(new)}$	2.9591	0.00700	422.7
Prev. unempl., $\beta_{unemp}^{\pi(new)}$	1.2326	0.00337	365.6

low-skilled jobs is 0.099 instead of 1 as in Copenhagen, the probability of ending up there would have been only 8.0 percent. The large regional differences in supply of jobs is therefore strongly reflected in the job probabilities.

Figure 15a also shows that the model can capture the share moving work location over the life cycle. There are some challenges of modelling the work transition probabilities for the younger individuals. Especially for those who have children where the model overpredicts the mobility while for low-skilled people it underprecits mobility. Looking specifically at the probability of moving work location to and from Copenhagen, Figure 16a illustrates that the share moving their job away from Copenhagen (out of all individuals in the data) shows a satisfactory fit though the model underpredits from age 40 onwards. The motivation for moving one's job conditional on the home location is shorter commute or higher wages. Commute distances are exogenous and thus independent of age, while wages have an age profile. We are aware wages may not exhibit enough variation across individuals as we do not allow for unobserved heterogeneity. Including this may improve on the fit since we would better capture whether the more mobile individuals are those who have a high unobserved fixed component of wages that they can bring with them when they move around. When zooming in on movers only in Figure 16b the fit is worse. The model cannot capture individuals who would like to move away from Copenhagen where there is such a high job

Figure 15: Model fit: share moving work location over the life cycle

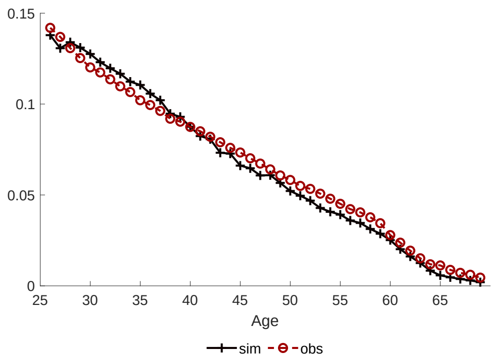

(a) Overall

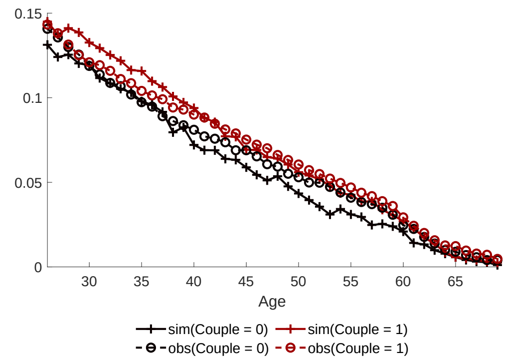

(b) By marital status

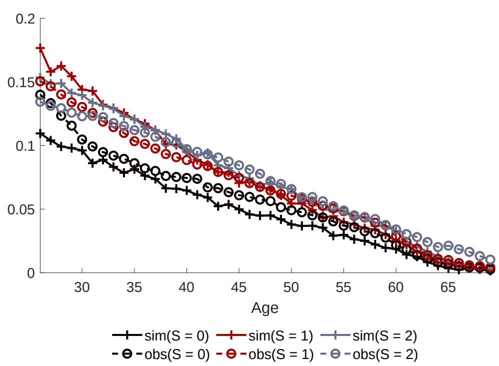

(c) By schooling

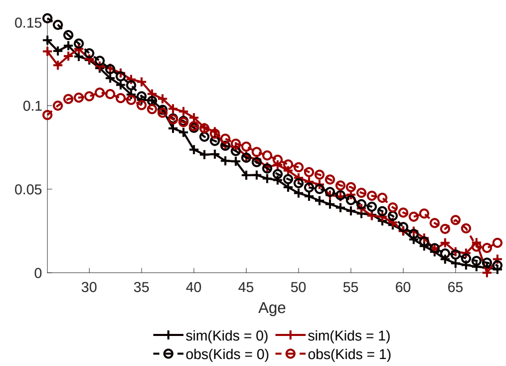

(d) By children

Figure 16: Model fit: share moving work location from and to Copenhagen over the life cycle

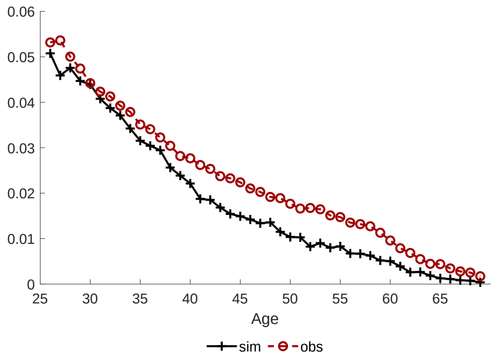

(a) From Copenhagen out of all

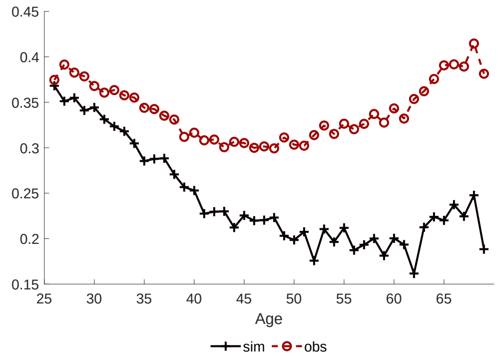

(b) From Copenhagen out of all movers

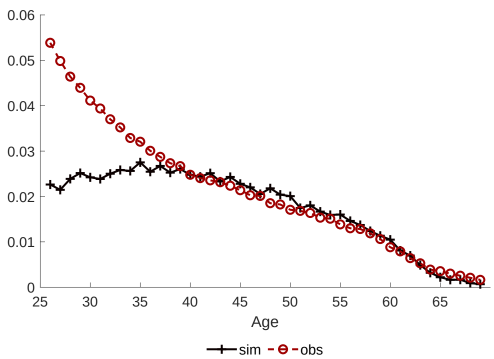

(c) To Copenhagen out of all

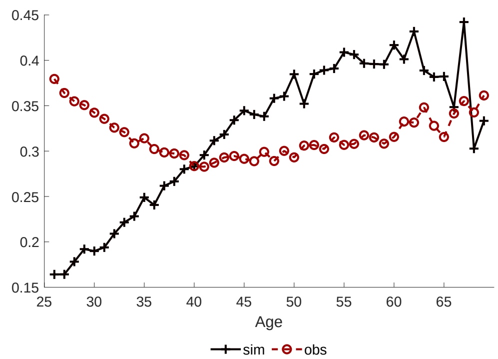

(d) From Copenhagen out of all movers

Note: Panels (a) and (c) show the share of all individuals in the data who move work location away from and to Copenhagen, respectively. Panel (b) and (d) show the share of all work location movers who move from and to Copenhagen, respectively.

density and thus high chance of being employed.

Considering instead the share working in Copenhagen in Figure 17a, the fit looks very good for the individuals older than 35. The heterogneity across individuals is also reflected in the model predictions. The work location decision is less well-captured for the young people because we do not model initial conditions or educational choice.

Looking at the share of individuals working in Copenhagen by their home municipality, the top panel of Figure 18 shows that the model captures the spatial distribution pretty well. It underpredicts the share somewhat for people also residing in Copenhagen. This can be improved by better capturing the share moving home location away from and to Copenhagen over the life cycle since that alone should make it more likely to also work in Copenhagen.

Figure 17: Model fit: share working in Copenhagen over the life cycle

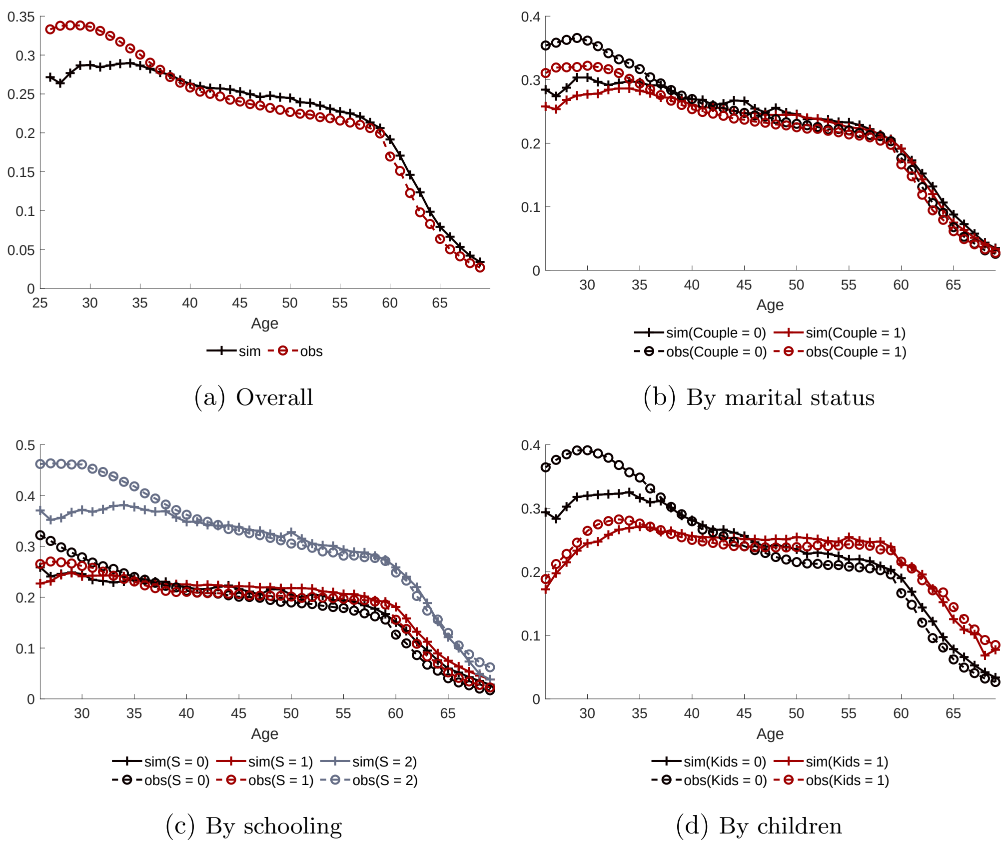

(a) Overall

(b) By marital status

(c) By schooling

(d) By children

Figure 18: Model fit: working in Copenhagen and commute times (hours)

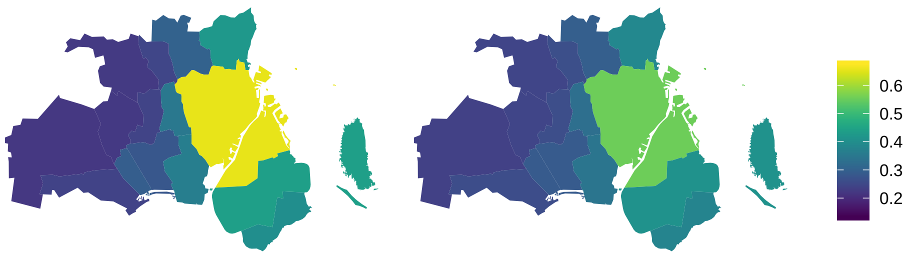

(a) Work in Cph. by home (obs.) (b) Work in Cph. by home (sim.)

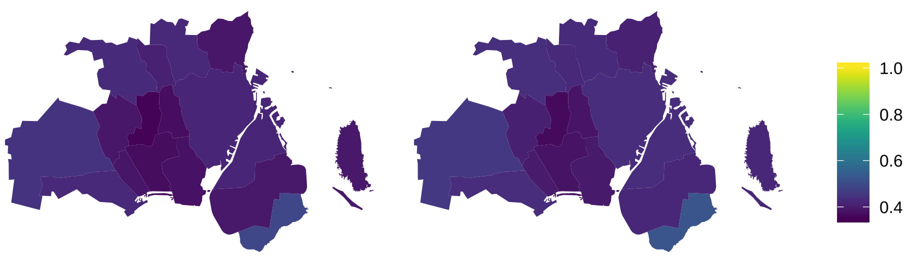

(c) Commute time by home (obs.) (d) Commute time by home (sim.)

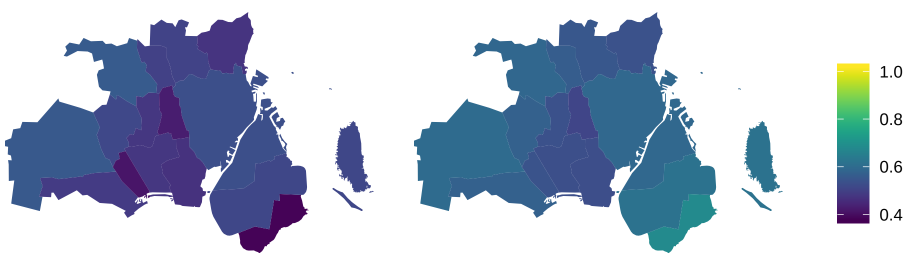

(e) Commute time by work (obs.) (f) Commute time by work (sim.)

Note: Panels (a) and (b) show the share of individuals in each home region who works in Copenhagen or Frederiksberg. Panels (c) and (d) show average commute time in hours by home region for employed individuals. Panels (e) and (f) show average commute time in hours by work region for employed individuals.

Table 6 provides estimates of the commute cost parameter, η_{ttime}, and disutility of working, c_{work}. The latter reflects the compensation one would require to take a job instead of being unemployed and corresponds to 280,393 DKK for a person with an annual (non-employment) income of 150,000 DKK. η_{ttime} indicates that an employed person with an income of 500,000 DKK would only be willing to commute one hour further is she earned an additional 42,685 DKK. The disutility of commuting in the data is therefore not overwhelming considering the fact that individual annual wage incomes typically are in the range 300,000-400,000 DKK.

Table 6: Commute Cost and Disutility of Work

	Coeff. Estimates	Standard Error	t-statistic
Cost of travel time, η_{ttime}	0.2369	0.00118	200.8
Disutil. of work, c_{work}	2.2163	0.00189	1175.6

As pictured in the middle and lower panels of Figure 18, the prediction error in the spatial allocation of commute times is low, especially by home locations. By work locations, the model predicts higher and a more uniform distribution of commute times than is observed in the data. The fact that average commute times are generally higher when splitting by work instead of home location is because people from Rest of Zealand also commute to the regions shown on the map. Figure 19a illustrates the commute time over the life cycle and across different types of individuals and it is predicted very accurately by the model. It is mainly for individuals above age 65 that the model starts to struggle, but there is also a strong selection among working individuals at that age. It is therefore not surprising that they cannot necessarily be compared to the younger working cohorts.

7.2 Baseline equilibrium

To assess whether the observed house prices in our data form an equilibrium at the housing market, we solve for the equilibrium prices following the procedure outlined in Section 6 and using the obtained structural estimates.

Figure 20 plots computed equilibrium prices against observed price data. The fit appears very good both in terms of the price ranking of the different regions as well as the overall levels. Here is it is important to emphasize that the model is estimated without explicitly

Figure 19: Model fit: commute time (hours) over the life cycle

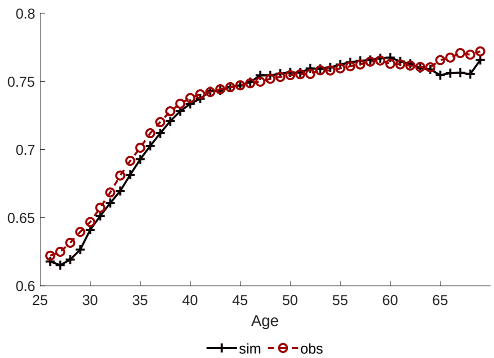

(a) Overall

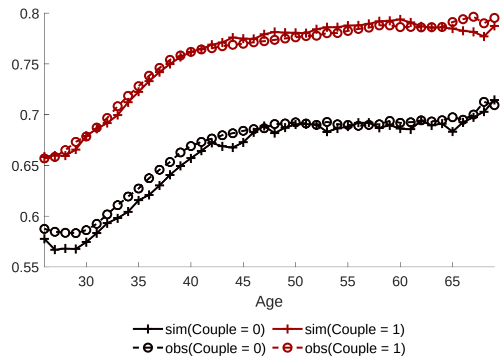

(b) By marital status

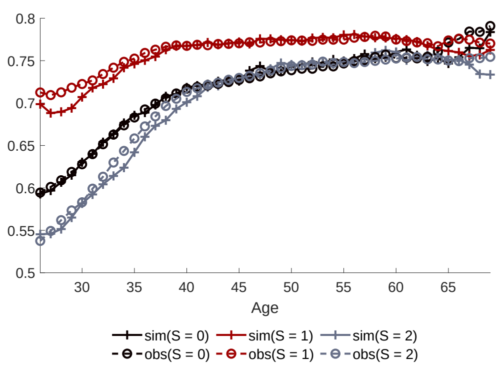

(c) By schooling

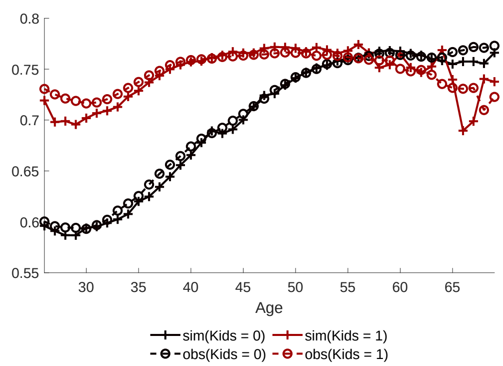

(d) By children

Figure 20: Relation between observed and baseline equilibrium house prices per sqm (100,000 DKK)

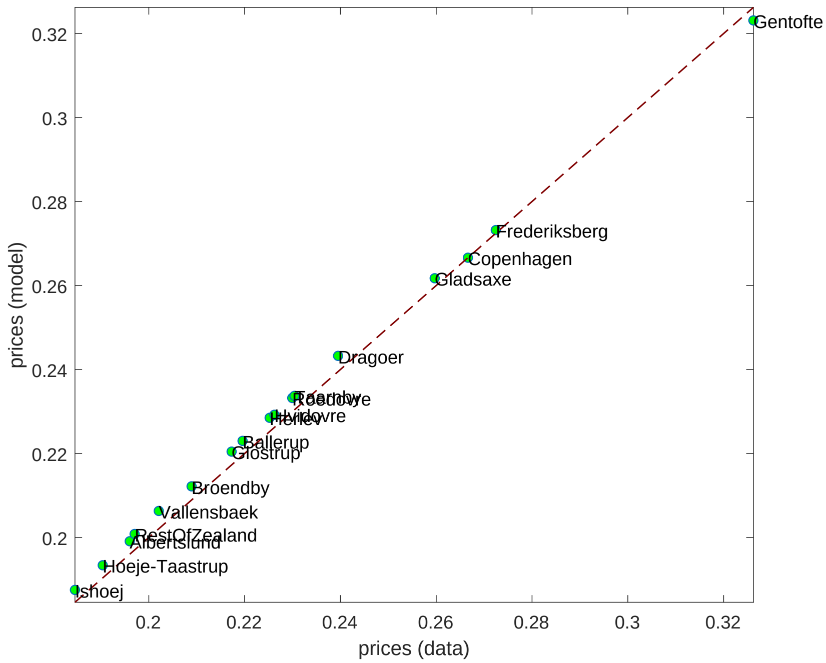

imposing that the housing market is in equilibrium. The fact that the equilibrium prices predicted from our estimated model closely track the observed house prices in the different regions provides a good in-sample validation of the many cross-equation restrictions implied by our modeling of location choices and demand for house size.

With the overall fit being exceptionally good, there is a slight overprediction of prices in the cheapest regions and an underprediction in Gentofte, the most expensive region. Our parsimonious modeling of individual income and the lack of savings are again among the potential explanations as to why the model does not fully capture why people are willing to pay such high prices in Gentofte.

7.3 Counterfactual equilibrium

In order to make a valid comparison between the baseline of the model and a counterfactual simulation, we account for the implied relocations and price changes that were to occur even in the absence of any policy change (due to demographic trends) by simulating the model forward a number of periods. In doing so, we obtain a simulated household-level panel

Figure 21: Structure of comparison between counterfactual and baseline

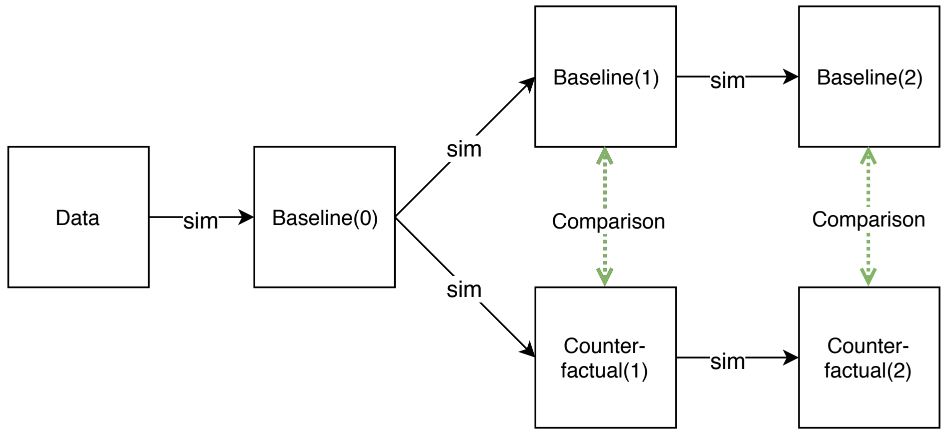

dataset with baseline outcomes. The baseline simulation starts at the empirical data on which the model was estimated. The outcome of simulating the model one period ahead from the empirical data yields the initial state for both the following baseline simulation steps as well as the counterfactual. This structure is illustrated in Figure 21, where the first baseline dataset is denoted Baseline(0). As Baseline(0) is the initial condition for the counterfactual, policy changes are imposed *at the beginning* of simulated period 1. At the end of period 1 it is therefore possible to identify all changes between baseline and counterfactual outcomes at the household level.

7.3.1 Counterfactual I: Increased housing supply

The first counterfactual experiment involves a 5 percent exogenous and permanent increase of the housing supply (square meters) in Copenhagen and Frederiksberg. Using the simulations for period 2, we study the implications for location choice, housing size, income sorting and equilibrium prices.

Table 7 summarizes the first four measures. As expected, the share living in Copenhagen and Frederiksberg increases, though just by 0.12 and 0.09 percentage points, respectively. Even though supply was constant in all other regions, the number of people living in Gentofte, Gladsaxe and Roedovre (which all share borders with Copenhagen) also show positive trends. As a result, the average housing size falls in these three regions. The demand for living in all the remaining regions drops, so in total the degree of centralization and urbanization has increased.

Individuals living in Copenhagen and Frederiksberg already can immediately adjust their housing demand upwards due to the lack of adjustment costs. However, some people from outside those regions, who did not prefer living there in baseline, are now inclined to relocate to those cities because there is a possibility of consuming more square meters. This starts the equilibrating process of some people moving out and substituting Copenhagen and Frederiksberg by Gentofte, Gladsaxe and Roedovre which are close substitutes in space. The increased demand for living in these nearby regions affects their equilibrium prices and prompts the original residents to consider moving too.

Looking at the resulting equilibrium prices, Figure 22 shows that all regions experience falling prices per square meter. This is especially true in Copenhagen, Frederiksberg, Gentofte and Roedovre despite the increased demand for living there. The average price per square meter in Copenhagen was 26,661 DKK and it falls by 750 DKK corresponding to 2.8 percent. Thus, locally in Copenhagen 56 percent of the supply shock is soaked up by falling prices. The lower prices in Copenhagen and Frederiksberg are caused by a more moderate increase in demand for living there than the increase in housing supply. In Gentofte and Roedovre there are two counteracting effects: increased demand for living there and spillovers from the generally lower price level in the GCA. The latter dominates. In the rest of the region the drop in prices is due to the substitution effect that induces people to move away and closer to the urban center.

Due to the reallocation of people across space, the sorting patterns have also changed in equilibrium. The second column of Table 7 shows the change in average income by home region and the third column the change in the within-region standard deviation in income. Copenhagen, Frederiksberg and Gentofte experience an increase in average income and incomes are more homoegeneous within those regions after the policy change, i.e. due to the resorting of individuals, the increased supply of housing does not induce many low-income households to live in Copenhagen and Frederiksberg in equilibrium despite the lower equilibrium prices. In Roedovre, on the other hand, incomes are lower and more heterogeneous. Dragoer stands out as the region with the highest increase in average income and the most

Table 7: Simulated changes in $t = 2$ in Counterfactual I

	Population Share % points	E(inc) %	Std(inc) %
Copenhagen	0.12	0.05	-0.05
Frederiksberg	0.09	0.13	-0.47
Ballerup	-0.01	0.15	0.11
Broendby	-0.02	-0.12	0.13
Dragoer	-0.05	0.94	-2.61
Gentofte	0.10	0.01	-0.46
Gladsaxe	0.06	-0.50	0.43
Glostrup	-0.01	-0.47	0.48
Herlev	-0.01	-0.01	0.08
Albertslund	0.00	0.16	0.00
Hvidovre	-0.01	-0.08	0.13
Hoeje-Taastrup	0.00	-0.11	0.21
Roedovre	0.02	-0.44	0.56
Ishoej	-0.02	0.09	-0.47
Taarnby	-0.03	0.00	-0.21
Vallensbaek	-0.01	0.08	-0.21
Rest Of Zealand	-0.20	-0.01	0.00

Note: Numbers are computed by subtracting baseline from counterfactual. Population share refers to the change in the share of all individuals who live in the region. E(inc) refers to the change in the average income of residents in the region. Std(inc) refers to the change in the standard deviation of income of residents in the region.

significant fall in the standard deviation of income. This is consistent with the idea that the lower-income original residents of Dragoer move towards Copenhagen, where they can now consume a satisfying number of square meters at a more reasonable price than in baseline. The average income of residents in Dragoer is indeed higher than in Copenhagen, so it is likely that the lower-income outmigrants from Dragoer to Copenhagen have an income above the average for the original Copenhagen residents.

In conclusion, the effect on commute time is negligible and the effects on residential location and sorting are more complex. We will investigate this in more detail in future work.

7.3.2 Counterfactual II: Increased commute costs

In the second counterfactual we increase the commute costs η_{ttime} by 50 percent. This might resemble an increase in monetary costs of traveling due to a removal of the mileage allowance (a tax benefit for commuters), lower subsidies on ticket prices, increased gasoline prices or alternatively increasing congestion in a uniform manner across all regions.

Figure 22: Simulated change in price per sqm in $t = 2$ in Counterfactual I (DKK)

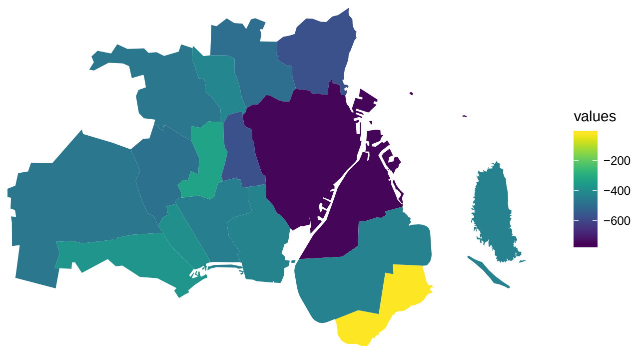

Note: Numbers are computed by subtracting baseline from counterfactual.

Figure 23: Distribution of simulated change in commute time in $t = 2$ in Counterfactual II (hours)

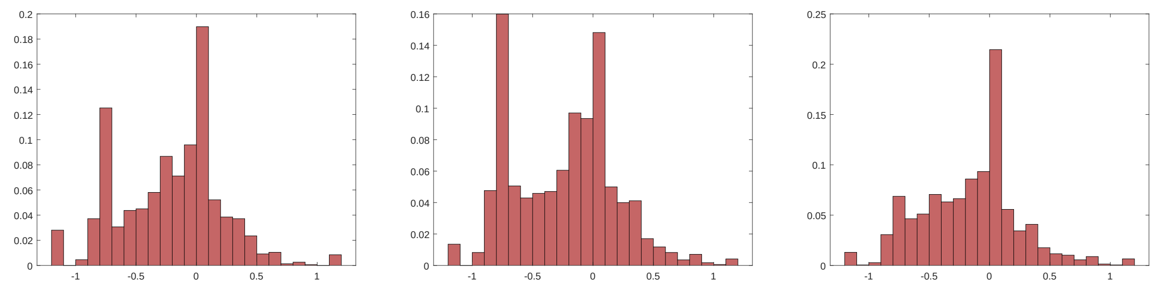

(a) Low education (b) Medium education (c) High education

Note: Changes in commute times are computed by subtracting baseline from counterfactual.

As commute costs rise, households will intuitively want to be closer to their jobs. This may be achieved through relocating their jobs, relocating home location or opting out of working entirely. The overall effect on commuting in the counterfactual is visualized by the histograms in Figure 23 showing differences in commute time when comparing with the baseline simulation. To be clear, for each individual in the counterfactual simulation we have subtracted their commute time in the baseline simulation. Further, we condition on individuals who have i) chosen another home location than in baseline and ii) were in employment in simulated period 0. We condition on i) and ii) for the sake of exposition. In particular, ii) avoids a large mass at 0 as a consequence of no commute time in non-employment coupled with the high persistence of the non-employment state.

Figure 23 clearly shows more mass in the negative support stemming from individuals lowering commute time in response to the increased costs. The figures also show a mass point around -0.75 hours for low- and medium-educated individuals. This is caused by individuals who in baseline lived in Rest of Zealand and worked in Copenhagen, yet after the change in commute costs decided to opt for non-employment. The same mass point does not occur for the highly educated, see Figure 23c, as they have much less incentive to opt out of the labor market due to higher opportunity costs of not working.

We would expect the relocations of individuals to be particularly evident in the most peripheral regions and this is confirmed in Table 8. The first two columns show the percentage of individuals who relocated their job away from each municipality in baseline and counterfactual. Most strikingly, individuals working in Rest of Zealand change their work location to a relatively large degree. This can be concluded as it only holds 8 percentage points more workers than Copenhagen, cf. Table E1 in appendix, while the relative increase in job relocations is much higher than for Copenhagen. Like Rest Of Zeland, the municipality of Ishoej is associated with long commute times due to its location on the southern perimeter of the GCA. Being on the perimeter, Ishoej and Rest of Zealand are relatively close substitutes in terms of work locations and we therefore also see significant increase in the relocations of jobs away from Ishoej.

As a case study of the model predictions, Table 9 displays detailed moving statistics for Ishoej. The first column shows the initial (simulated period 0) residential locations of individuals working in Ishoej. Predominantly, workers of Ishoej lived in Rest of Zealand and many therefore had long commute times. In the counterfactual state of higher commute costs, those workers would be particularly discouraged from continuing to work. We see this pattern indirectly in column two, which displays the distribution of work locations for those who switched work location between periods 0 and 2. 24.5 percent of people employed in Ishoej in period 0 did not work in period 2. As the third column shows, 66.8 percent of these non-employed people lived in Rest of Zealand in period 0, hence underlining the discouraging effect of higher commute costs. Note also that Gentofte, which is located the furthest away from Ishoej, displayed the second-highest share of non-employed residents in period 2.

Table 8: Simulated share of relocations of work (wl) and home (rl) and price change in $t = 2$ for Counterfactual II

	Baseline ($t = 2$) wl	Counterfactual ($t = 2$) wl	Baseline ($t = 2$) rl	Counterfactual ($t = 2$) rl	Δprice (DKK)
Copenhagen	3.66	3.98	1.15	1.00	59.93
Frederiksberg	6.51	6.28	1.98	1.99	86.24
Ballerup	7.21	7.84	3.44	3.08	73.96
Broendby	10.34	11.04	3.46	2.94	77.54
Dragoer	27.24	26.26	4.78	5.57	57.43
Gentofte	7.53	7.46	2.80	2.71	6.62
Gladsaxe	7.72	7.84	1.57	1.17	81.80
Glostrup	11.22	11.27	3.56	3.74	247.17
Herlev	11.94	11.98	2.72	2.94	248.75
Albertslund	12.99	13.30	2.37	1.75	51.89
Hvidovre	9.74	10.05	2.09	1.78	108.70
Hoeje-Taastrup	9.94	10.83	2.28	1.97	13.45
Roedovre	13.97	14.69	4.27	3.90	1.70
Ishoej	20.93	22.58	5.80	5.25	17.12
Taarnby	11.97	12.22	5.96	5.99	96.89
Vallensbaek	36.17	36.34	7.29	7.49	136.72
Rest of Zealand	0.87	1.41	0.30	0.45	-148.13
Non-employment	5.78	5.69	-	-	-

Column two also indicates which work regions are the closer substitutes to Ishoej. The municipalities of Taarnby, Roedovre and Taastrup attract the most workers from Ishoej, although Vallensbaek and Albertslund are closer to Ishoej than both Roedovre and Taarnby. However, these are relatively small labor markets so the probability of getting a new job prohibits workers from relocating there.

Returning to Table 8, column five shows the change in equilibrium prices in period 2 between counterfactual and baseline. We note that all regions except Rest of Zealand experience slightly increasing prices, while the prices in Rest of Zealand declines. This is a direct consequence of lower demand for residing in Rest of Zealand and a substitution towards the GCA where commute times are lower. Correspondingly, column three and four show the share of outmigrants from each home region in baseline and counterfactual. There is a net inflow to the CGA as the propensity to move away from Rest of Zealand is higher in the counterfactual while it is lower in Copenhagen. These two regions dominate the picture due to their sizes, see Table E1 in appendix.

Table 9: Simulated distribution of locations in $t = 2$ for $t = 0$ workers in Ishoej in Counterfactual II

	Home region of workers in $t = 0$ (%)	New wl of job movers in $t = 2$ (%)	Home region of job movers when new $wl = \emptyset$ in $t = 2$ (%)
Copenhagen	16.8	7.3	0.6
Frederiksberg	3.8	0.7	2.6
Ballerup	1.4	0.5	0.6
Broendby	2.0	1.1	-
Dragoer	0.5	0.2	0.3
Gentofte	4.6	0.9	15.3
Gladsaxe	2.1	0.5	1.9
Glostrup	0.9	0.3	1.6
Herlev	0.9	0.5	1.3
Albertslund	1.0	2.1	1.3
Hvidovre	2.4	4.1	1.9
Hoeje-Taastrup	3.7	15.7	2.2
Roedovre	1.3	14.9	1.0
Ishoej	10.2	-	0.6
Taarnby	0.9	20.0	0.6
Vallensbaek	1.2	3.5	1.3
Rest of Zealand	46.4	3.0	66.8
Non-employment	-	24.5	-

8 Conclusion

In this paper we developed and empirically estimated a structural dynamic equilibrium model of joint home and work location decisions for individuals and estimated it using Danish administrative data. We found that overall the empirical fit of the model is very good. We focused on the Greater Copenhagen Area (GCA) and analyzed the counterfactual effects of i) increasing the housing supply in Copenhagen and Frederiksberg by 5 percent and ii) increasing commute costs by 50 percent.

We found that the increase in housing supply resulted in relocations towards Copenhagen and Frederiksberg such that the degree of urbanization increased. The relocations did not completely offset the increase in the housing supply, so the average housing size also increased in those two regions. In total, the equilibrium prices dropped in all regions and especially in Copenhagen, where they fell by 2.8 percent. The sorting of individuals was also affected. Hence, Copenhagen and Frederiksberg were characterized by richer and more homoegeneous households on average after the policy change.

The increase in commute costs not only caused an anticipated relocation leading to a decrease in average commute times, but also to a significant labor supply effect. In particular, a significant share of residents of Zealand outside of the Copenhagen region who worked in Copenhagen ended up in non-employment in the counterfactual. This effect is more pronounced for low- and medium-educated workers. The downward movement of the equilibrium prices follows the decline in labor participations, which is in line with the higher incentive to live within the GCA where commute times were lower.

Overall, the model developed and estimated in the paper provides valuable insights into our understanding of the location and movement patterns among Danish households, which are driven by the cost of living and commuting, and are very heterogeneous in the population. The current implementation of the model is not free of strong simplifying assumptions, but even in their presence it proves to be a very valuable tool, capable of explaining important variation in the data, and enabling us to undertake interesting counterfactual experiments.

Among most significant limitations of the current implementation are the effective disregard of the time dimension of the data (especially in the dimension of developing amenities in different regions), and the static equilibrium house price calculations. The regions can be less aggregated, and a wider area of the country rather than the GCA can be used for estimation. Inclusion of the equilibrium wage settlement into the consideration is another obvious dimension for improvement. Even under the assumption of short term dynamics in the labor market similar to the housing market (so that the supply of jobs is constant) the wages can be treated similarly to house prices and be determined in the spatial equilibrium. All of these improvements, although requiring additional work and computational time, are straightforward to implement. Even though we do acknowledge all the limitations and relevant extensions mentioned above we leave the implementations for future research.

References

Ahlfeldt, G. M. and McMillen, D. P. (2015). The vertical city: The price of land and the height of buildings in chicago 1870-2010. SERC discussion papers, London School of Economics.

Ahlfeldt, G. M., Redding, S. J., Sturm, D. M., and Wolf, N. (2015). The economics of density: Evidence from the berlin wall. *Econometrica*, 83(6):2127–2189.

Albouy, D. and Stuart, B. (2019). Urban Population and Amenities: The Neoclassical Model of Location. *Forthcoming International Economic Review*.

Alonso, W. et al. (1964). *Location and Land Use. toward a General Theory of Land Rent.* Publications of the Joint Center for Urban Studies of the Massachusetts Institute of Technology and Harvard University. Cambridge: Harvard University Press.

Attanasio, O. P., Bottazzi, R., Low, H. W., Nesheim, L., and Wakefield, M. (2012). Modelling the demand for housing over the life cycle. *Review of Economic Dynamics*, 15(1):1–18.

Bayer, P., Keohane, N., and Timmins, C. (2009). Migration and hedonic valuation: The case of air quality. *Journal of Environmental Economics and Management*, 58(1):1–14.

Bayer, P., McMillan, R., Murphy, A., and Timmins, C. (2016). A dynamic model of demand for houses and neighborhoods. *Econometrica*, 84(3):893–942.

Bishop, K. C. (2012). A dynamic model of location choice and hedonic valuation. *Unpublished, Washington University in St. Louis.*

Borjas, G. J. (2000). The economic progress of immigrants. In *Issues in the Economics of Immigration*, pages 15–49. National Bureau of Economic Research Conference Report. Chicago: University of Chicago Press.

Buchinsky, M., Gotlibovski, C., and Lifshitz, O. (2014). Residential Location, Work Location, and Labor Market Outcomes of Immigrants in Israel. *Econometrica*, 82(3):995–1054.

Dahl, G. B. (2002). Mobility and the return to education: Testing a roy model with multiple markets. *Econometrica*, 70(6):2367–2420.

Favilukis, J., Mabille, P., and Van Nieuwerburgh, S. (2019). Affordable housing and city welfare. Technical report, National Bureau of Economic Research.

Gallin, J. H. (2004). Net migration and state labor market dynamics. *Journal of Labor Economics*, 22(1):1–21.

Gould, E. D. (2007). Cities, workers, and wages: A structural analysis of the urban wage premium. *The Review of Economic Studies*, 74(2):477–506.

Guglielminetti, E., Lalive, R., Ruh, P., and Wasmer, E. (2017). Home sweet home? job search with commuting and unemployment insurance. Working paper.

Halket, J., Nesheim, L., and Oswald, F. (2015). The housing stock, housing prices, and user costs: The roles of location, structure and unobserved quality. Technical report, Cemmap working paper, Centre for Microdata Methods and Practice.

Hicks, J. (1932). The theory of wages. *The Economic Journal*, 43(171):460–472.

Holt, F. M. (1997). *Family Migration Decisions: A Dynamic Analysis*. PhD thesis, University of Virginia.

Kennan, J. and Walker, J. R. (2011). The effect of expected income on individual migration decisions. *Econometrica*, 79(1):211–251.

Kuminoff, N. V., Smith, K. V., and Timmins, C. (2013). The new economics of equilibrium sorting and policy evaluation using housing markets. *Journal of Economic Literature*, 51(4):1007–62.

Mangum, K. (2015). Cities and labor market dynamics. *Andrew Young School of Policy Studies Research Paper Series*, (15-05).

Manning, A. and Petrongolo, B. (2017). How local are labor markets? evidence from a spatial job search model. *American Economic Review*, 107(10):2877–2907.

McFadden, D. (1974). The measurement of urban travel demand. *Journal of public economics*, 3(4):303–328.

McFadden, D. (1978). Modeling the choice of residential location. *Transportation Research Record*, (673).

McMillen, D. P. (2006). Testing for monocentricity. *A Companion to Urban Economics. Oxford: Blackwell*, pages 128–140.

Mills, E. S. (1967). An aggregative model of resource allocation in a metropolitan area. *The American Economic Review*, 57(2):197–210.

Muth, R. F. (1969). *Cities and Housing: the Spatial Pattern of Urban Residential Land Use*. Studies in business & society / Univ. of Chicago. Graduate School of Business. Chicago University Press.

Nathanson, C. G. (2019). Trickle-down housing economics. Working paper.

Ogawa, H. and Fujita, M. (1980). Equilibrium land use patterns in a nonmonocentric city. *Journal of regional science*, 20(4):455–475.

Oswald, F. (2019). The effect of homeownership on the option value of regional migration. *forthcoming Quantitative Economics*.

Ransom, T. (2016). The effect of business cycle fluctuations on migration decisions. *Available at SSRN 2741117*.

Roback, J. (1982). Wages, rents, and the quality of life. *Journal of political Economy*, 90(6):1257–1278.

Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of political economy*, 82(1):34–55.

Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica: Journal of the Econometric Society*, pages 999–1033.

Rust, J. (1988). Maximum likelihood estimation of discrete control processes. *SIAM Journal on Control and Optimization*, 26(5):1006–1024.

Schmutz, B. and Sidibé, M. (2015). Job search and migration in a system of cities. *Economic Research Initiatives at Duke (ERID) Working Paper*, (181).

Sjaastad, L. A. (1962). The costs and returns of human migration. *Journal of political Economy*, 70(5, Part 2):80–93.

Tiebout, C. M. (1956). A pure theory of local expenditures. *Journal of political economy*, 64(5):416–424.

Topel, R. H. (1986). Local labor markets. *Journal of Political economy*, 94(3, Part 2):S111–S143.

Tsivanidis, N. (2019). Evaluating the impact of urban transit infrastructure: Evidence from bogota's transmilenio. Working paper.

Tunali, I. (2000). Rationality of migration. *International Economic Review*, 41(4):893–920.

A Geographic Classifications

Table A1: Overview of geographical classifications in Denmark

Danish	English	# units	Comment
Danmark	Denmark	1	
Regioner	Regions (states)	5	
Landsdele	Provinces	11	10 excl the island Bornholm
Amter	Counties	16	No longer exists
Valgkredse	Constituencies	92	
Kommuner	Municipalities	98	Reform in 2007: from 271 mun. to 98.
Trafikzoner	Traffic/LTM zones	907	Defined by DTU's traffic model
Sogne	Parish	2,201	

Figure A1: Municipalities of Denmark and urbanized areas

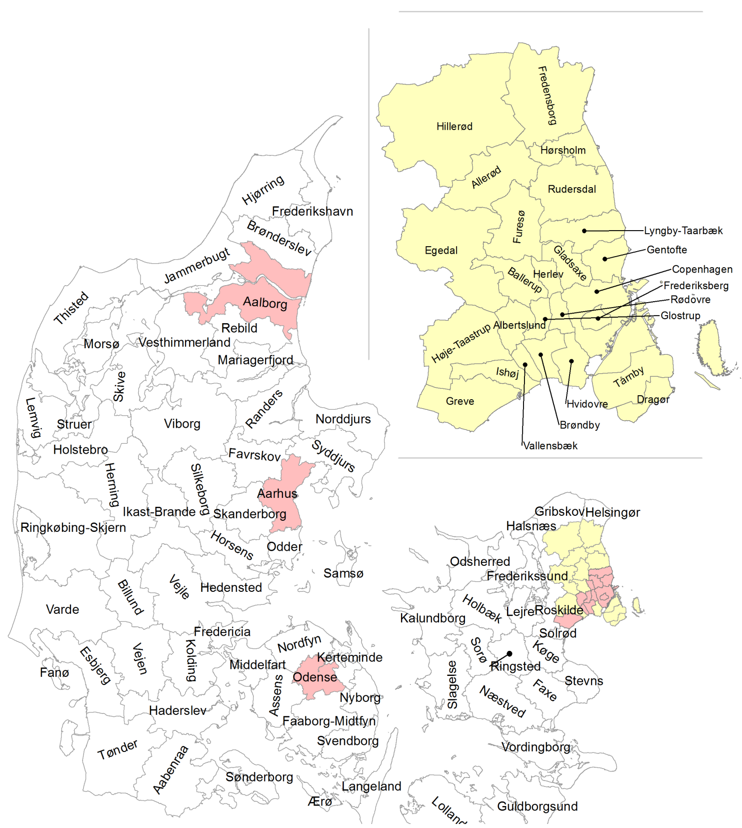

Note: Pink areas indicate an urban municipality. The yellow area corresponds to the main part of the greater Copenhagen area.

Figure A2: Provinces and main islands

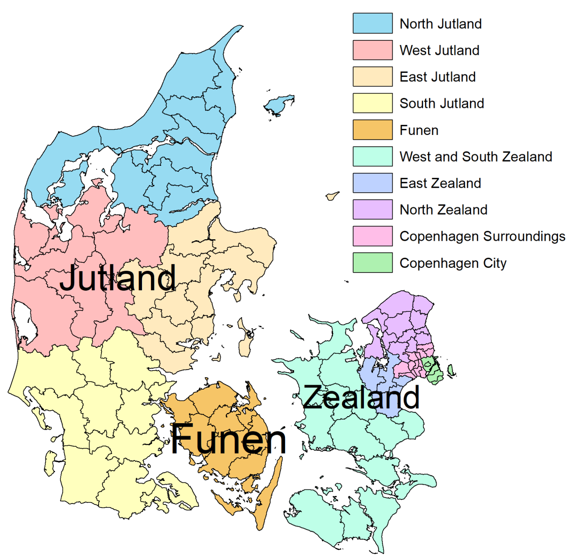

B Amenities

Table B1: Summary statistics of job density by work region

Region	Low educ.	Medium educ.	High educ.
Copenhagen	1.0000	1.0000	1.0000
Frederiksberg	0.1117	0.1197	0.0966
Ballerup	0.1036	0.1517	0.0943
Broendby	0.0747	0.0976	0.0501
Dragoer	0.0074	0.0102	0.0064
Gentofte	0.0847	0.0962	0.1020
Gladsaxe	0.0915	0.1255	0.1042
Glostrup	0.0593	0.0886	0.0547
Herlev	0.0468	0.0721	0.0524
Albertslund	0.0639	0.0945	0.0463
Hvidovre	0.0783	0.1016	0.0605
Hoeje-Taastrup	0.0990	0.1351	0.0579
Roedovre	0.0474	0.0714	0.0310
Ishoej	0.0258	0.0332	0.0149
Taarnby	0.1011	0.1062	0.0353
Vallensbaek	0.0110	0.0147	0.0082
Rest Of Zealand	0.0508	0.0722	0.0397
Funen	0.0613	0.0893	0.0473
Jutland	0.0857	0.1211	0.0617

Note: Job density is defined as the number of jobs by education group and work region normalized by the value in Copenhagen. The numbers have been averaged over time.

C Wage Regressions

Table C1: Estimates from OLS of Log Real Wages for Low-Skilled Workers by Region

Work Region	age	age^2	$\mathbb{I}_{rw_{t-1}=\emptyset}$	$Constant$	R^2	N
Copenhagen	0.1587 (0.000)	-0.0017 (0.000)	-0.7765 (0.000)	9.0397 (0.000)	0.2332	410758
Frederiksberg	0.1554 (0.000)	-0.0016 (0.000)	-0.7561 (0.000)	9.0395 (0.000)	0.2327	43752
Ballerup	0.1313 (0.000)	-0.0014 (0.000)	-0.7811 (0.000)	9.8425 (0.000)	0.1976	47008
Broendby	0.1036 (0.000)	-0.0011 (0.000)	-0.7962 (0.000)	10.4932 (0.000)	0.1605	35442
Dragoer	0.1321 (0.000)	-0.0015 (0.000)	-0.6105 (0.000)	9.6267 (0.000)	0.1873	3412
Gentofte	0.1281 (0.000)	-0.0013 (0.000)	-0.7183 (0.000)	9.6782 (0.000)	0.1787	39708
Gladsaxe	0.1601 (0.000)	-0.0018 (0.000)	-0.8380 (0.000)	9.1874 (0.000)	0.2683	38132
Glostrup	0.1349 (0.000)	-0.0014 (0.000)	-0.6912 (0.000)	9.7162 (0.000)	0.2223	23925
Herlev	0.1306 (0.000)	-0.0014 (0.000)	-0.7055 (0.000)	9.7836 (0.000)	0.1958	19150
Albertslund	0.0957 (0.000)	-0.0010 (0.000)	-0.7419 (0.000)	10.5897 (0.000)	0.1527	28222
Hvidovre	0.1169 (0.000)	-0.0013 (0.000)	-0.6683 (0.000)	10.0812 (0.000)	0.1801	35960
Hoeje-Taastrup	0.1184 (0.000)	-0.0013 (0.000)	-0.7218 (0.000)	10.1158 (0.000)	0.1919	41393
Roedovre	0.1160 (0.000)	-0.0013 (0.000)	-0.6916 (0.000)	10.0640 (0.000)	0.1855	19878
Ishoej	0.1164 (0.000)	-0.0013 (0.000)	-0.7522 (0.000)	10.0493 (0.000)	0.1774	11720
Taarnby	0.1456 (0.000)	-0.0016 (0.000)	-0.7374 (0.000)	9.6359 (0.000)	0.2093	47816
Vallensbaek	0.1315 (0.000)	-0.0014 (0.000)	-0.7503 (0.000)	9.7778 (0.000)	0.2039	4907
Rest Of Zealand	0.1486 (0.000)	-0.0016 (0.000)	-0.6826 (0.000)	9.3810 (0.000)	0.2174	76475

Note: Standard errors in parentheses. $\mathbb{I}_{rw_{t-1}=\emptyset}$ means unemployed in $t-1$.

Table C2: Estimates from OLS of Log Real Wages for Medium-Skilled Workers by Region

Work Region	age	age^2	$\mathbb{I}_{rw_{t-1}=\emptyset}$	$Constant$	R^2	N
Copenhagen	0.1141 (0.000)	-0.0013 (0.000)	-0.7975 (0.000)	10.3419 (0.000)	0.1654	524053
Frederiksberg	0.1157 (0.000)	-0.0013 (0.000)	-0.8044 (0.000)	10.2659 (0.000)	0.1784	58739
Ballerup	0.1062 (0.000)	-0.0012 (0.000)	-0.7160 (0.000)	10.5949 (0.000)	0.1504	90715
Broendby	0.0961 (0.000)	-0.0011 (0.000)	-0.7564 (0.000)	10.8255 (0.000)	0.1517	53930
Dragoer	0.0925 (0.000)	-0.0011 (0.000)	-0.7435 (0.000)	10.7246 (0.000)	0.1376	5442
Gentofte	0.1021 (0.000)	-0.0012 (0.000)	-0.7419 (0.000)	10.6416 (0.000)	0.1561	59172
Gladsaxe	0.1247 (0.000)	-0.0014 (0.000)	-0.7785 (0.000)	10.2222 (0.000)	0.2258	69497
Glostrup	0.0985 (0.000)	-0.0011 (0.000)	-0.7235 (0.000)	10.7475 (0.000)	0.1556	45200
Herlev	0.0927 (0.000)	-0.0010 (0.000)	-0.6233 (0.000)	10.8365 (0.000)	0.1359	38145
Albertslund	0.0862 (0.000)	-0.0010 (0.000)	-0.7069 (0.000)	11.0337 (0.000)	0.1417	50102
Hvidovre	0.0894 (0.000)	-0.0010 (0.000)	-0.6742 (0.000)	10.8940 (0.000)	0.1324	56422
Hoeje-Taastrup	0.0905 (0.000)	-0.0010 (0.000)	-0.7153 (0.000)	10.8998 (0.000)	0.1334	71102
Roedovre	0.0911 (0.000)	-0.0010 (0.000)	-0.7055 (0.000)	10.8677 (0.000)	0.1451	36186
Ishoej	0.0979 (0.000)	-0.0011 (0.000)	-0.6555 (0.000)	10.7029 (0.000)	0.1548	17606
Taarnby	0.0900 (0.000)	-0.0010 (0.000)	-0.6377 (0.000)	10.9085 (0.000)	0.1303	57936
Vallensbaek	0.1047 (0.000)	-0.0012 (0.000)	-0.7684 (0.000)	10.5610 (0.000)	0.1578	8112
Rest Of Zealand	0.0974 (0.000)	-0.0011 (0.000)	-0.6999 (0.000)	10.7593 (0.000)	0.1682	102252

Note: Standard errors in parentheses. $\mathbb{I}_{rw_{t-1}=\emptyset}$ means unemployed in $t-1$.

Table C3: Estimates from OLS of Log Real Wages for High-Skilled Workers by Region

Work Region	age	age^2	$\mathbb{I}_{rw_{t-1}=\emptyset}$	$Constant$	R^2	N
Copenhagen	0.1687 (0.000)	-0.0018 (0.000)	-0.6917 (0.000)	9.1518 (0.000)	0.2132	657976
Frederiksberg	0.1634 (0.000)	-0.0017 (0.000)	-0.7167 (0.000)	9.1166 (0.000)	0.2159	71351
Ballerup	0.1444 (0.000)	-0.0015 (0.000)	-0.6027 (0.000)	9.8524 (0.000)	0.1779	62758
Broendby	0.1444 (0.000)	-0.0015 (0.000)	-0.6122 (0.000)	9.7670 (0.000)	0.1700	27438
Dragoer	0.1398 (0.000)	-0.0015 (0.000)	-0.7592 (0.000)	9.6609 (0.000)	0.1923	3559
Gentofte	0.1401 (0.000)	-0.0015 (0.000)	-0.7552 (0.000)	9.8928 (0.000)	0.1556	77232
Gladsaxe	0.1504 (0.000)	-0.0016 (0.000)	-0.6465 (0.000)	9.7113 (0.000)	0.1837	64861
Glostrup	0.1264 (0.000)	-0.0013 (0.000)	-0.6007 (0.000)	10.1571 (0.000)	0.1662	33987
Herlev	0.1094 (0.000)	-0.0011 (0.000)	-0.6071 (0.000)	10.4400 (0.000)	0.1397	31325
Albertslund	0.1337 (0.000)	-0.0014 (0.000)	-0.7433 (0.000)	10.0164 (0.000)	0.1529	21939
Hvidovre	0.1121 (0.000)	-0.0012 (0.000)	-0.6212 (0.000)	10.3636 (0.000)	0.1456	35407
Hoeje-Taastrup	0.1411 (0.000)	-0.0015 (0.000)	-0.6668 (0.000)	9.8275 (0.000)	0.1724	32240
Roedovre	0.1186 (0.000)	-0.0012 (0.000)	-0.6439 (0.000)	10.1954 (0.000)	0.1365	15335
Ishoej	0.1238 (0.000)	-0.0013 (0.000)	-0.5989 (0.000)	10.0317 (0.000)	0.1495	8132
Taarnby	0.1382 (0.000)	-0.0014 (0.000)	-0.7239 (0.000)	9.7441 (0.000)	0.1756	19408
Vallensbaek	0.1295 (0.000)	-0.0014 (0.000)	-0.6547 (0.000)	10.0063 (0.000)	0.1642	4764
Rest Of Zealand	0.1467 (0.000)	-0.0015 (0.000)	-0.5958 (0.000)	9.6913 (0.000)	0.2091	131409

Note: Standard errors in parentheses. $\mathbb{I}_{rw_{t-1}=\emptyset}$ means unemployed in $t-1$.

D Structural Estimates

Table D1: Regional Amenities

	Coeff. Estimates	Standard Error	Z-statistic
α_{rl} (1)	0.0153	0.00051	30.1
α_{rl} (2)	-0.9733	0.00145	-673.3
α_{rl} (3)	-1.2263	0.00178	-690.8
α_{rl} (4)	-0.6359	0.00268	-237.3
α_{rl} (5)	0.7848	0.00134	583.6
α_{rl} (6)	-0.2120	0.00085	-249.2
α_{rl} (7)	-1.0813	0.00196	-550.6
α_{rl} (8)	-0.8991	0.00178	-504.3
α_{rl} (9)	-1.5117	0.00217	-695.8
α_{rl} (10)	-0.8425	0.00128	-657.0
α_{rl} (11)	-1.5901	0.00196	-811.1
α_{rl} (12)	-0.7930	0.00138	-576.3
α_{rl} (13)	-1.7885	0.00252	-708.8
α_{rl} (14)	-0.7490	0.00136	-551.9
α_{rl} (15)	-1.4207	0.00249	-571.4
α_{rl} (16)	-1.1823	0.00159	-743.5

E Counterfactual

Table E1: Share of individuals in each region in baseline $t = 0$ (pct)

	Baseline: wl	Baseline: rl
Copenhagen	20.1	19.5
Frederiksberg	2.3	4.1
Ballerup	2.9	1.7
Broendby	1.8	1.3
Dragoer	0.4	0.6
Gentofte	2.4	4.1
Gladsaxe	2.5	2.5
Glostrup	1.7	0.7
Herlev	1.4	0.9
Albertslund	1.6	0.9
Hvidovre	1.9	1.8
Hoeje-Taastrup	2.2	1.8
Roedovre	1.2	1.3
Ishoej	0.8	0.7
Taarnby	1.9	1.5
Vallensbaek	0.6	0.6
RestOfZealand	28.1	55.9
Non-employment	26.2	-

F Description of the Data Sources

This section provides details of how the sample we use is constructed from individual Danish registers.

F.1 Individual background characteristics

The population register BEF is posted on January 1st each year and lists all individuals who have their officially registered address in Denmark. Each individual in the register is represented by an anonymized version of their official social security number called PNR. PNR is used to merge BEF with other registers with individual-specific data. From BEF we know the age and gender of the individual, whether she has children, how old the children are, and if she lives with a partner (married or not). UDDAUPD is informative of the highest completed education of the individual on a very detailed level down to the field of study. By using a table from Statistics Denmark that translates finer codes into broader categories we can reduce the number of education categories to the three categories we use in the estimation of the model: low (no more than high school), medium (vocational or short-length further education) and high education (bachelor degree or more). The education register is updated every October. To make sure observations from BEF and UDDAUPD are as close in time as possible, we merge UDDAUPD in year t on to BEF from year $t + 1$ via PNR.

F.2 Addresses and home moves

Importantly, BEF also contains an anonymized version of the individual's home address and an unmasked code for the home municipality, parish and other administrative geographic regions. In 2007, a municipality reform took place in Denmark which reduced the number of municipalities from 271 to 98 municipalities. This caused a change in the home addresses in the register (as they are only unique within a municipality), but we have used a key file from Statistics Denmark that translates old addresses to their new version post 2007. We therefore use the definition of municipalties from 2007.

For each individual we also know when they moved into the address they are currently at. Since our model is formulated at an annual basis we define the individual to live at the address where she lived during most of the year. If she moves to a new address in e.g. May and stays there for the rest of the year, this end-of-period address will be her home region choice, but if she did not move until August, we would record her beginning-of-period address as her home choice.

F.3 Labor market information

Data on workplace and other workplace-related variables such as industry and occupation come from the Integrated Database for Labor Market Research (IDA). IDA consists of different panels: one for personal data on employees (IDAP), another one for employments (IDAN) and one for workplaces (IDAS). We mainly use IDAN which has a record for every combination of individual, employment and year. The information about an individual's employment in IDAN comes from the Central Tax Information Sheet Register (Centrale Oplysningsregister) until 2008 and from eIncome (also located at the tax authorities) for the remaining period.

An individual can have several employments during a year. The register is posted by the end of November each year and groups individuals' employments into either employed wage-earners, employer (A), self-employed (S) or co-working spouse (M). All groups are mutually exclusive. The group of wage-earners is then further divided into main occupation (H), sideline occupation or another November occupation or most important non-November occupation (the two latter categories only available from 2004 onwards). For each individual who has more than one wage-earner job in November, we use information from the main occupation (H) which is determined by the largest source of income. Individuals who do not work are either classified as unemployed or outside the labor force. We classify individuals as unemployed if they according to the register are coded as being on leave (including parental leave and sick leave), unemployed by the end of November, participating in unemployment activation (short-term jobs finacially supported by the public sector) or on rehabilitation. This information comes from IDAP. We define people to be outside the labor force whenever an individual is not unemployed and is recorded as being outside the labor force, studying, retiree, early retiree, on transitional allowance or on social security benefits. IDAP also has a variable showing for how many days the person has been registered as (un)employed during the year and we use this as the individual's (un)employment rate.

Each workplace has an anonymized version of its address which is recorded in IDAN. This anonymous address can be linked to non-masked municipalities, parishes or traffic zones (LTM zones). In some cases employments in IDAN cannot be assigned to a registered workplace. Instead, Statistics Denmark assigns the home address as the workplace (a so-called fictious workplace). This is typically the case for employees who work from home or at several workplaces, e.g. cleaners or community nurses.

To model characteristics of the work regions, we compute a job density measure for each region. We define this measure as the number of jobs within three levels of education and normalize by the corresponding level in Copenhagen. The number of employees by region is available from IDAN and after merging with BEF we know the education group too.

Since IDAN and IDAP are posted in November, we merge IDAN and IDAP year t on to BEF year $t + 1$ via PNR.

F.4 Income

Data on income is available from INDUPD for a given year but not split on different employments within the year[14]. We are able to distinguish between total income, wage income and transfer income before and after taxes though and use income measures before tax. For people whom we classified as working in November, we use their annual wage income divided by their employment rate in the year according to IDAP. For people who are unemployed in November we use their transfer income divided by their unemployment rate and for those outside the labor force we use their total income.

[14] There exists another register BFL which has wage income for each combination of employment, individual and month, but only since 2008.

F.5 Commute time

Commute time data come from The Danish Traffic Model (LTM) which has been developed by researchers at The Technical University of Denmark (DTU). They have divided Denmark into 907 traffic zones (LTM zones) and modelled commute time between each pair of regions. They model commute time for different transport modes (car, public or walk/bike) and exploit information on the road network, speed limits, congestion, bus and train timetables including waiting times, and bike paths. The traffic model has been run for 2002 and 2010 using the road network and public transport schedules for each year. Since our model is formulated in terms of municipalties, not LTM zones, we compute a commute time measure by each transport mode between any pair of LTM zones within a municipality pair. For a given pair of LTM zones in a municipality pair, we use the commute time from the mode with the shortest commute time. We then weigh the commute time of each observed LTM pair in the municipality pair with its estimated number of trips by that mode from the traffic model and thereby get a trip-weigted average commute time between any pair of municipalities. The difference between the 2002 and 2010 simulations of the traffic model is due to changes in the road network or bus and train plans. Commute time by walking or biking is constant.

From LTM we also get data on travel distances between each zone. We do the same exercise for distance as we did for travel time to get an average distance measure between each pair of municipalities.

F.6 Property prices, home ownership and home characteristics

Information on property prices come from the sales transaction register EJSA. EJSA holds a record of each sale in Denmark including sales price, type of sale (e.g. single-family house or commercial), number of square meters sold, and type of post-sale ownership (e.g. private or business). We deflate all sales prices by the consumer price index with 2011 as the reference year. We use only private sales and disregard properties with commercial-only purpose. On top of that we clean the data on sales prices using the same criteria that Statistics Denmark uses for official statistics on property prices, i.e. property value must exceed the lot value[15],

[15] Property and land value measures come from the register EJVK.

the property must not have been sold more than once on the same day and is sold on open market terms.

Data on home ownership come from the EJER register. It links every housing unit in Denmark with a PNR of the owner and define an owner as someone who owns more than zero percent of the property. In order to link EJER and EJSA, we exploit the unique housing unit identifier which is available in both registers. This enables us to merge EJSA with our household panel via PNR. Since EJSA is posted on January 1st and EJER on October 1st each year we merge EJSA year t with EJER year $t-1$ and then merge EJSA year t with the household panel of year t.

For home characteristics we use BOL which is based on BBR (The Central Register for Buildings and Dwellings). It is a register with a record of each property in Denmark and gives information on characteristics such as the number of rooms, bathrooms and most importantly square meter for living space. BOL in year t can be merged to the houseold panel in year t via the housing unit identifier.

Chapter 3

Altruism and Tax Incentives: Parental Landlords in Denmark

This paper investigates how parents in Denmark have invested in apartments for the purpose of providing rentals for their adult children, primarily while enrolled in further education. The investments are favored by a tax system that offers significant benefits and by surging urban housing prices. Parents are therefore in a position to support their adult children while making a sound investment. Using full-population register data, I document that both altruistic and financial motives play an important role for parents' decision to invest in rental apartments. I complement the analysis with a collective model of the investment decision to consider the effects of existing tax benefits on welfare and housing prices. The model shows that the tax system increases inequality in welfare among adult children and tends to push up housing prices.

This project has received funding from Realdania and the Danish Research institute for Economic Analysis and Modeling which is gratefully acknowledged.

1 Introduction

Affordable rentals have been in notoriously short supply in the major Danish cities for the past decades. This has hurt the housing prospects of low-income households, students and young professionals who are barred from the freehold market due to liquidity and income constraints. Because of the lacking supply of affordable housing in cities, landlord arrangements within families have gained popularity in Denmark since the 1990s. That is, parents to young adults have purchased apartments in the major Danish cities for the purpose of renting them to their offspring as a way to alleviate a tight housing situation, especially while studying.

Borrowing from the UK glossary, these arrangements are dubbed parental buy-to-let investments (or BTLs for short). Buy-to-let investments are indubitably spurred by altruism of parents towards their children. There are, however, additional incentives at play to direct the decision process. For one, the Danish tax system features important benefits for parents who engage in rental investments. Secondly, the surge in urban housing prices throughout the 2000s (and 2010s) created large gains on real estate investments which parents were able to share with their children by selling them the apartments at a discount.

In this paper, I investigate empirically and theoretically how parents respond to the incentives for making buy-to-let investments. This investment behavior is interesting because it involves elements of both inter-vivo giving, portfolio management and tax-planning, and the incentives underlying each element are not necessarily aligned.

Inter vivo-giving and the allocation of wealth between (living) generations is important for the study of inequality at large and has received attention in recent years. Unfortunately, comprehensive data sources of monetary transfers between family members are scarce as they require either access to bank statements or questionnaires on the subject. Bank statements have only on rare occasions been made available for academic studies while surveys historically have been confined to asking only for crude assessments of intra-family transfers. A contribution of this paper is therefore to create a full population register based data set that is revealing about parental economic support through buy-to-let investments in

Denmark between 1994 and 2015.

A set of regressions is performed on this data to interpret behavior on i) the extensive margin of whether to provide a buy-to-let or not, ii) the intensive margin concerning the value of the buy-to-let conditional on investing and iii) the rebate offered to children in case they buy it from their parents. Overall, the regressions document a behavior of parents tending to equalize wealth and consumption among family members. This is seen from the facts that lowering child income is associated with higher probability of making buy-to-let investments, reduced size and hence rent of the buy-to-let, and increased rebate offered to those children who buy the apartments from their parents.

The second contribution of the paper is to formulate a collective model of the buy-to-let investment which encompasses the relevant tax regulation to clarify the mechanisms at work. The model runs over two periods from the perspective of a (partly) altruistic parent who cares about both the wellbeing of the child and her own consumption.[1] The parent faces a discrete-continuous choice of investing in a buy-to-let apartment (discrete), the apartment's size (continuous), the down payment (continuous) and how much of the rent income is saved up for later taxation (continuous). Conditional on no investment, the child must find a rental on his own. Conditional on investment, the child lives in the apartment in both periods but buys the apartment in the beginning of period two at a rebate decided by the parent (equal to the down payment).

The family gains on the buy-to-let arrangement because the child's rent payments now flow to the parent rather than to a third party, because the child can be offered a lower rent than obtained in the private rental market and because wealth can be transferred untaxed through the rebate on the apartment.

The model is solved to yield the following predictions:

- *The probability of making buy-to-let investments is increasing in the difference between income of a parent and child. This is mainly because a parent can compensate the child*

[1] From here on, the parent is 'her' and the child is 'he'.

for the lower income by charging a lower rent pr m^2 than otherwise obtained in the rental market.

- *The rebate offered to children when selling the apartment is increasing in the income difference between parent and child, as this is another way for the parent to equalize welfare within the family.*
- *The distribution of parents' pre-tax income will show bunching around the top-tax bracket because of an incentive to shift income between periods.*

The predictions are confirmed by the regressions presented in the first part of the paper. Although the predictions are tested in a reduced form, the setup readily lends itself to a full structural estimation in future research.

The solution to the model also shows that buy-to-let investments and the associated tax regulation work to increase inequality in housing outcomes among children. Furthermore, the model is extended to an equilibrium setup where the price pr m^2 on the market for children's housing is endogenous. In this stylized scenario, the buy-to-let arrangement causes the price to increase by 34%.

One of the richest studies of inter-vivo giving in terms of background variables and timespan is found in McGarry (2016) who analyzes the U.S. Health and Retirement Survey. This survey includes questions on cash transfers from family members over a 20-year horizon and documents that parents do compensate children for lower life-time earnings. It also shows that transfers from parents correlate with home purchases of children. However, the questions regarding family cash transfers only demand a crude estimate in asking whether any transfers above USD 1,000 were received. Similarly, the study by Altonji et al. (1997) looks at the 1988 wave of the PSID where respondents are asked to recollect the magnitude of any *loan or gift* above USD 100 (1987) received from parents during the preceding year. Using this measure, they implement an econometric test of parental altruism which rejects that the parents' transfers are consistent with a model of parental altruism.

For comparison, the transfers involved in the rebates on buy-to-let apartments are measured on a continuous scale (they may even become negative, i.e. apartments are overpriced) and the sample size is considerably larger, which gives a more nuanced characterization. The regressions show that parents standing to loose money on the buy-to-let investment largely roll back rebates, even though their income is higher.

A recent paper by Andersen et al. (2020) has, as a novelty in the literature, succeeded in obtaining and combining bank records with background register data for customers at the largest Danish retail bank. They show that parents (in contrast to siblings and friends) are the main providers of informal insurance against adverse income shocks and financial distress. They also show that the replacement rate of lost income is highest at the lowest child income deciles, consistent with a notion of parental altruism. Furthermore, the replacement rate is strongly dependent on the parental income level.

An interesting approach to circumvent the problem of poorly observed transfers is found in Boar (2019), who instead consider parents' savings behavior. She shows that parents increase their savings if children experience heightened income risk, implying that they are bracing their finances for making future transfers. In a closely related paper, Kaplan (2012) finds that young people are often allowed to move back home as compensation for negative labor market shocks which helps explain low precautionary savings on their part.

A worry one might have is that transfers from parents to children appear as gifts yet are in reality loans to alleviate temporary credit constraints. Hochguertel and Ohlsson (2009) entertain this possibility when also looking at the Health and Retirement Survey and conclude that observed inter-vivo transfers to children must be compensatory gifts rather than loans. Based on the buy-to-let data, one cannot rule out that children might be expected to make repayments later in life due to the rebates or decreased rent. However, it is in Denmark perfectly legal for family members to extend loans to each other with the intermediation of a bank, which appears to be an easier solution all together if the parents are merely lending money out.

A point the current paper is that parents are responsive to tax-induced incentives in their decisions to make within family transfers because these can normally be planned well ahead of time. Empirically, Bernheim et al. (2004) and Joulfaian (2004) provide strong evidence for this notion. They show that the timing of bequests in the U.S. is very sensitive to changes in taxation, implying that parents are strategic in their transfers of wealth. The same result prevails in Escobar and Ohlsson (2019); they study a change in the Swedish inheritance tax code where recipients of bequests were able to avoid taxation by transferring the tax liable bequests to their own children. The effects of this incentive is clearly pronounced in their data through massive bunching around the first kink in the inheritance tax schedule. Finally, Sommer (2017) studies the effects of the progressive German inheritance tax (which includes taxation on gifts) and show that people are much more responsive to taxes levied on inter-vivo gifts compared to inheritance.

A burgeoning literature, starting with Bernheim et al. (1985), is modeling the decision process of transfers and inheritance within the family as involving an exchange motive rather than being pure gift giving. For instance, Barczyk et al. (2019) demonstrate empirically that bequests can be a tool for parents to ensure their children's commitment to provide care in old age. A similar finding is provided by Brown (2006). The analysis of buy-to-let investments in the current paper is related to this strand of literature in that financial gains accruing to parents are an important determinant of the investment.

The paper proceeds as follows. Section 2 clarifies the institutional setup surrounding buy-to-let investments and the incentives on behalf of parents. Section 3 describes data sources and provides descriptive statistics of the population and Section 4 presents the regressions on buy-to-let investment behavior. Section 5 formulates the model of buy-to-let investments, compares its predictions with the empirical findings and provides comparative statics. Section 6 concludes the paper.

2 Institutional setup

2.1 Tax incentives for buy-to-let investors

This section introduces the tax rules and legal framework around buy-to-let investments and the incentives these produce. Descriptive evidence on the impact of incentives is presented alongside the facts concerning the institutional setup. We may note from the outset, that the primary incentives for buy-to-let investments in the tax system are that the interest payments are partly deductible, that rent income can be shifted between periods to minimize taxation and that wealth can be passed on to the children untaxed. The following paragraphs explain how each of these benefits are built into the tax system while additional motives are discussed in the next subsection.

The legislation governing buy-to-let investments is such that parents must set up an enterprise with a bank account that is detached from their private economy and keep a detailed financial record. At the end of each fiscal year, parents can choose freely which of three tax schemes this enterprise is subject to. These are i) *personal income* taxation ii) *capital returns* taxation and iii) *firm* taxation. It is in general easier to manage the capital returns scheme compared to the firm taxation scheme, which requires assistance from an accountant, but also typically less financially advantageous. The personal income scheme is very simple to manage, but also financially inferior in most cases.

Table 1 shows the number of annual observations of parent-households with buy-to-let apartments within each tax scheme. Note that a parent household is counted under a given tax-scheme if just one parent files under that scheme. The reasoning here is that a parent household will be able to utilize the benefits of that scheme if just one member files under it. However, a few households are filing under both schemes, likely because they are running additional enterprises. The table also shows the median of pr parent income[2] for households using each scheme. The firm taxation scheme is the relevant option for most buy-to-let

[2] See note under Table 1.

Table 1: Parents with buy-to-let investments in each tax scheme. Observations pooled across years.

Tax scheme	Buy-to-let parents 1994-2015	Median income 2010-2015
Personal income	34,613	395,623
Capital return	36,478	428,064
Firm taxation	80,564	496,082
Firm and capital return	1,704	462,817

Notes: First column displays the number of parent-child observations with a buy-to-let apartment pooled across years. Second column is median taxable income pr. adult in buy-to-let families, DKK 2015 prices. For discretionary purposes, the measure displayed is the average of all incomes at the 50th percentile. These numbers are based the years 2010-2015 only for a clear illustration.

owners while personal income and capital return taxation are roughly equally popular. There is also a clear gradient in household income with respect to the chosen scheme, with firm taxation being chosen by the highest earning parents and personal income taxation chosen by the lowest. Why this is the case will be clear in the following exposition of the tax system.

To help explain the taxation of buy-to-let investors, Table 2 presents the key components of the Danish income tax system which are relevant for the context. The rates and cut-offs are based on 2009, yet the basic structure has remained in place since 1994 (with most cut-offs and tax rates varying over time). However, in 2010 the middle income tax was repealed and the deduction of positive net capital income, t^{cap}, was increased to 40,000. The rate on the bottom government tax has increased steadily since 2010 and is now 12.1%.

The *personal income* scheme considers the buy-to-let apartment to be a part of one's private economy. Rent income from the buy-to-let factors into the *personal income* measure in table 2, while the interest payments on the accompanying mortgage (negative capital income) factors into *taxable income*. Therefore, the interest payments are deducted from health and local government taxation, but they do not affect the basis for government taxation unless one already has positive net capital income.[3]

The *capital returns* scheme requires a weak separation of one's private economy and the

[3] In Denmark, households most commonly have negative net capital income due to mortgage payments.

Table 2: Danish income taxation in 2009

Income measures		
Personal income		{Salary + profits from own firm + public benefits + pension payouts + fringe benefits} × (1 - labor tax [8%])
Taxable income		Personal income + net capital income - deductions
Tax	**Rate**	**Base**
Gvt tax, bottom	∼5%	{Personal income + net capital income > 0} > t^{bottom}
Gvt tax, middle	6%	{Personal income + net capital income > 0} > t^{middle}
Gvt tax, top	15%	{Personal income + net capital income > t^{cap}} > t^{top}
Local gvt tax	∼25%	Taxable income
Health tax	8%	Taxable income
Cut-offs (DKK1000s)		
t^{bottom}	41	
t^{middle}	280	
t^{top}	336	
t^{cap}	0	

Notes: Married couples are allowed to transfer capital income between spouses. Positive net capital income of one spouse can therefore avoid taxation to the extent that the other spouse files negative net capital income (or capital income < t^{cap} wrt. to the top tax base).

buy-to-let apartment. It allows owners to calculate an imputed capital income from the property which is deducted from rent income.[4] The remainder is then filed as personal income and the imputed capital income figures into taxable income. Due to interest payments, the owner will also have negative capital income to counterweight the positive imputed capital income the *taxable income* measure. In this way, income from the buy-to-let is moved away from the more heavily taxed personal income measure to taxable income, which is only basis for local taxation.[5]

Finally, the *firm taxation* scheme presupposes that the buy-to-let economy is entirely separate from private accounts. One significant advantage of the firm taxation scheme is that only *net profits*, income minus interest payments, are tax liable as personal income when transferred to the private economy. Hence, interest payments are fully deductible.[6] The other significant advantage is the *business cycle adjustment* option. The firm taxation scheme is originally intended for the self-employed whose income is highly volatile and therefore punished by the progressive income taxation. Buy-to-let owners using the scheme are therefore allowed to retain profits in the buy-to-let enterprise for later withdrawal and taxation. The business cycle adjustment option is also partially available under the capital returns scheme, where 25% of annual profits can be retained. This is pertinent for parents to children in higher education, as they know that their personal income will fall in the coming years due to retirement.

[4]Rent income is net of maintenance expenses.

[5]Historically, the rate used to impute capital income has been on par with effective mortgage rates for standard loans.

[6]Maintenance expenses are likewise deducted.

Figure 1: Income bunching among buy-to-let owners compared to others.

(a) Density of taxable income in 2012. (b) Density of taxable income around top kink, 2000-2015.

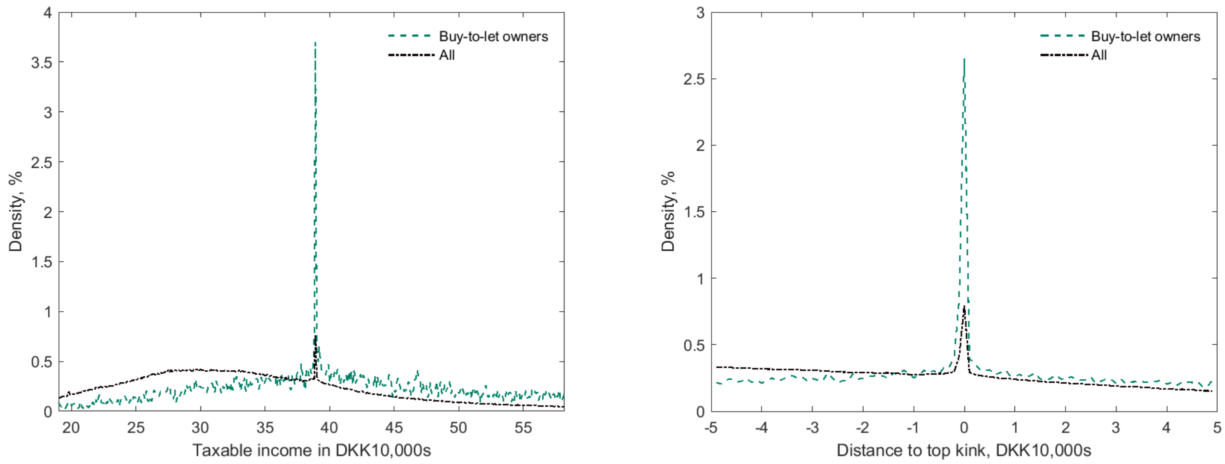

Notes: Incomes are in 2015-DKK when pooled across years. The top tax kink was located at 389,000 in 2012.

This feature of the tax system empirically entails bunching in taxable income just around the top tax cut-off among self-employed workers, le Maire and Schjerning (2013). The same is therefore expected for buy-to-let owners, and is indeed observed to be the case. Figure 1 Panel (a) displays bunching in taxable income around the top tax cut-off for a specific year, 2012, while Panel (b) aggregates taxable income distributions for the years 2000-2015 in a neighborhood of annual top tax cut-offs. We observe excessive bunching in taxable income among buy-to-let owners throughout the period, although the tendency is stronger in later years.

Another important tax incentive for prospective buy-to-let investors is the option of legally selling the buy-to-let apartment to the child at 85% of its public valuation. This has lately proven to be an effective vehicle for circumventing gift and inheritance taxes.[7] Public valuations have generally relied on rather conservative estimates and were updated every two years. Yet they will remain frozen for at least the period 2012-2021 due to technical issues of implementing an automated system. Public valuations have therefore in effect been kept at

[7]In 2020, the value of gifts above 67,100 DKK should be filed as personal income. Inheritance tax rate is 15%.

financial crash levels throughout the recent urban housing boom, making room for significant untaxed transfers between parents and children.

2.2 Further motives for buy-to-let investments

In addition to incentives stemming from tax regulation, parents may see a buy-to-let apartment as a lucrative investment opportunity as well as a means to support their children. Starting with the investment motive, note that housing prices in Copenhagen (incl. Frederiksberg) have surged over the past 30 years, see Figure 2. The other major university cities, Aarhus and Odense, also experience rising prices during the mid 2000s, but the recovery after the crisis has been slower.

Part of the reason for the high costs of housing may be the strict regulation imposed on the rental sector which effectively works to decrease the supply of rentals. One element in the regulation is that a landlord may not issue a temporary lease contract unless the landlord is planning to move into the home at the end of the contract. As a consequence, an investor in real estate who only wants to hold on to an apartment for, say, 5 years, will have to sell the apartment at year 5 as a rental rather than as a freehold. That is, unless the renter has moved out of his own accord.

Figure 2: Price pr m^2 for apartments in largest university cities.

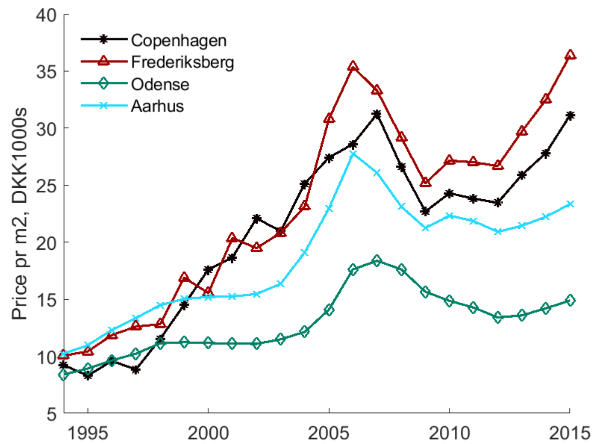

Notes: Prices are in DKK 2015 and based on median sales prices pr m^2 for freehold apartments traded on the open market. Frederiksberg is a separate municipality within Copenhagen.

Large property companies have little interest in detached rental units so the market for single rental units is thin, making small scale investments in rentals unappetizing. However, renting out to one's children while they are studying is an effective way of picking a renter with whom informal agreements about the length of stay can be made. Therefore, parents with a positive outlook on prices will have a unique opportunity of betting on the housing market, which would otherwise have been too cumbersome and risky.

Finally, investing in a buy-to-let apartment presents an opportunity for parental altruism towards the children, in addition to the inheritance and gift motives, as student housing and cheap rentals have been in excess demand for the entire period. The applicants on waiting lists reported by student housing companies have typically measured 2-3 times the number of housing units in stock[8], while access to social housing is even more restricted. Parents offering an affordable rental to children who have not lucked out in the housing lottery are therefore extending a genuine relief of stress. This shows up in the survey by Hjelmar (2009), where children in buy-to-let apartments express the highest degree of satisfaction with their

[8]Numbers pulled from a report on public housing programs for the young by Municipality of Copenhagen (2019).

housing situation while those in rented rooms and temporary leases are at the bottom.

A related issue is that parents can support their children economically by charging a cheap rent on the buy-to-let. Officially, a rent reduction received by the child from the parent is taxable, so in theory parents should charge a market based rent. To the extend that parents do wish to cut the rent, they have considerable room to do so before it catches the tax authority's eyes. This is due to the opacity of what rent should be required for an apartment that has no previous rent history, is situated in a regular housing block, and lives on a market where existing rent contracts vary extensively because of regulation.[9]

Unfortunately, there are no official records of rent charged by landlords in Denmark, but it is possible to impute it for renters who receive rent relief subsidies. Yet at the current state of the paper, I have not been able to this for children in buy-to-let apartments (who will often qualify for the subsidy), but it will be part of future work. An indication that parents are typically generous with the rent is found in Hjelmar (2009), where respondents state the share of income spent on housing. Buy-to-let tenants spend a smaller share of their income on housing than young people in any other form of housing; for example, 65% of buy-to-let tenants spend less than a quarter of monthly income while this only goes for 20% in a rented room.

3 Data and descriptives statistics

Data sources and structure

The data available for the empirical analysis is high quality register data covering the full population of Danish households between 1994-2015, obtained from Statistics Denmark. I define a parental buy-to-let observation as child who lives in a dwelling[10] which is owned by at least one biological parent and both parents must reside at another street address.[11]

[9] Parents can also consult websites like www.foraeldrekoebsguiden.dk which sells legal advice on, inter alia, how low rents may be set without raising red flags.

[10] Apartment or in rarer cases a house.

[11] I.e. a different OPGIKOM.

Identifying parental buy-to-let cases therefore requires merging address data[12] with property ownership data[13]. Ownership is recorded as a percentage in the registers so that both parents may be recorded as owners. The same registers are used to identify other real estate held by parents.

I consider the relevant family entity for the analysis to be child-parents pairs. Parents are assumed to act as a unitary decision maker vis-a-vis a child. Acting as a unitary decision maker implies that the income of parents is pooled, that only the highest level of education is recorded etc. Each record of the final data thus consists of pooled parental data and individual child data. In case the biological parents of a child are divorced, a record is created for each of the two new households, and the respective income of biological parents is pooled with their new spouses', if one is recorded, income.[14] A child with divorced parents will therefore occur twice within a year. If only one of the biological parents in a divorced couple is registered as the owner of a buy-to-let apartment, the child will not figure as a buy-to-let case in the record of the non-owner parent.

The income measure of parents should optimally consist of all earnings except incoming rent payments derived from the buy-to-let to avoid conflating outcomes with the driving forces. Parental income is therefore defined as salaries, public transfers, private pension payments and fringe benefits[15] in so far parents are not self-employed. If the parent is self-employed, one cannot distinguish rental income from regular business income and I therefore include all self-employed earnings[16] in the income measure. To be clear, I define a parent as self-employed if they file income under one of the self-employed taxation schemes and if their major source of income is self-employed business. The income of children is merely all taxable income consistent with the *personal income* measure in Table 2.[17] Pooled income of 2-person parental

[12]The register BEF

[13]The register EJER

[14]Note that a spouse in the definitions of Statistics Denmark is either a person one is married to, or a person of different gender with whom one shares address.

[15]From the register IND

[16]NETOVSKUD_13

[17]PERINDKIALT_13

Table 3: Characteristics of families with children aged 18-30, 1994-2015.

	Investment in buy-to-let	
	No	**Yes**
Income, children (100K)	1.957	1.795
Income pr. parent in parental unit (100K)	3.387	5.251
Parental units w. income above top tax cut-off (%)	32.8	64.0
Child in short cycle education (%)	48.5	30.3
Child in long cycle education (%)	33.1	68.3
At least one parent with high education (%)	33.5	63.6
Number of siblings	2.6	2.4
Biological parents not living together (%)	53.9	34.8
Properties owned by parent	0.89	1.65
Price pr m^2 in parent area (100K)	0.127	0.152
Cross-sectional units in total	2,398,698	
Average observations pr. unit	10.2	
Region	**Buy-to-let cases**	
Copenhagen	31,137	
Aarhus	11,725	
Aalborg	3,618	
Odense	2,792	
Esbjerg	509	
Rest of Denmark	8,115	

Notes: Incomes and prices are in DKK 2015-prices. Number of siblings includes children from non-biological parents, ie. half-brothers and half-sisters. Properties owned by parents includes all real estate and summer cottages.

units is rescaled by 2 to get the income pr parent. The income of children with spouses is not pooled. To incorporate the effects of the progressive tax system, I calculate the total amount of income liable to top taxes within a parental unit. This cannot be observed directly in data, and I therefore implement the tax simulator developed in le Maire and Schjerning (2013).

Finally, any ongoing education of children is obtained together with the highest level of completed education for both children and parents.[18]

[18]The register UDD.

Descriptives

Table 3 describes data associated with the cohorts of children born in the period 1970-1995 for the years 1994-2015 while they are between 18-30 years of age. The upper part of the table is split on whether a child lives in a parental buy-to-let at any point in the sample period. Accordingly, the education statistics and marital status of parents are given by their max over the period.

It is clear that children living in buy-to-lets come from privileged homes; income of parents is 50% higher, the parents are better educated and divorce rates are lower. Children in buy-to-let apartments have an almost twice the propensity to take long cycle education. Interestingly, parents investing in buy-to-let apartments are also much more prone to invest in other properties. The number of properties owned by parents investing in buy-to-let apartments is subtracted the buy-to-let apartment, so that most of them can be assumed to own their own home, the buy-to-let and a third property (a summer cottage for instance).

As explained above, investing in a buy-to-let is particularly beneficial if one is above the top tax bracket. The third entry of table 3 therefore features the simulated share of parental units in which at least one parents is above the cut-off for the top tax bracket. This share is about twice as large for buy-to-let parents.

Figure 3: Age profile of children in buy-to-let apartments, 1994-2015.

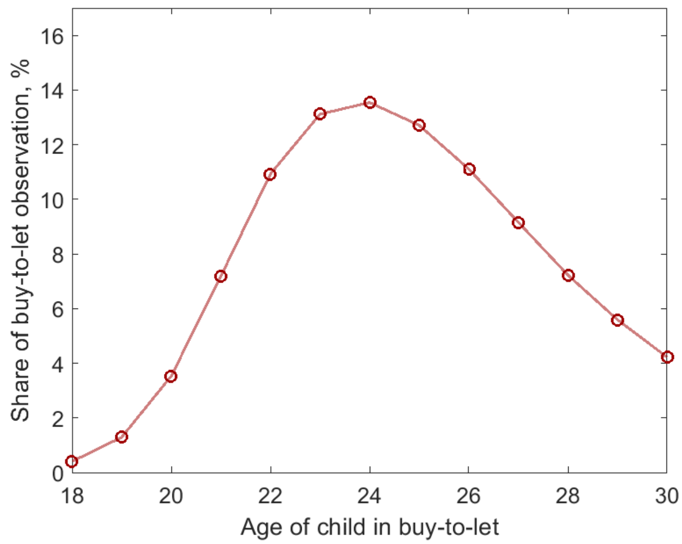

Figure 3 shows the age profile of buy-to-let tenants. The majority of buy-to-let arrangements are made when the child is between 20 and 24 years of age, which coincides well with enrollment in higher education.[19] As the children get older and graduate they start moving out, although a handful are still holding on to the apartment when they turn 30.

4 Empirical analysis of buy-to-let investments

In this section, I describe the empirical relationships between family characteristics and buy-to-let investment behavior. For the discrete choice on the extensive margin of investing in a buy-to-let or not, I implement a standard dynamic random effects probit model. The results are consistent with parents providing economic support for their children when investing in buy-to-let apartments while reaping tax benefits. The same is true on the intensive margin, which concerns the value of the value of chosen apartments. Child income shows to be an important determinant for the value of the buy-to-let, implying that their preferences are being accommodated.

[19] Note that Danish university students have historically been late starters with one or two sabbatical years between high school and university being the norm in the 1990s and 2000s.

I also consider the behavior concerning selling the buy-to-let; specifically I test whether parents who sell the apartment to their children are providing a discount compared to the market value and if this discount depends on relative incomes. This is indeed the case, suggesting that the intuitions from a model of altruistic parental behavior obtain in data.

4.1 Extensive margin of buy-to-let investments

To investigate the patterns in data regarding the choice on the extensive margin of investing in a buy-to-let, I apply the dynamic random effects probit model of Wooldridge (2005) with a modification suggested in Rabe-Hesketh and Skrondal (2013). It is necessary to apply a dynamic model to the buy-to-let decision because there are large fixed costs in obtaining and departing with an appartment. Such fixed costs induce strong state dependence which must be accounted for.

Furthermore, it is important to handle the problem of initial conditions problem of dynamic models. The initial conditions problem shows up in the available data on families since a sizeable fraction of children in the relevant age group experience changes in the composition of the parental household; i.e. divorces, remarrying and death of parents. It is doubtful that the random effect in the buy-to-let decision is unchanged when parents divorce or remarry, especially because it introduces new family members, so cross-sectional units with such events must either be dropped or split into subperiods with stable family composition. To avoid such sample selection, I choose to split cross-sectional units with changing family composition into stable sub-units. The implication is that although I observe practically every child from age 18 and on, some panels will start at later points in life.

Another problem is that including periods in which the child has not considered moving away from home at all will bias the coefficients on other explanatory variables, as the outcome will always be zero for those periods. Because there is not one age, at which everyone desires to move away from home, it is unclear when the econometrician should let the panel begin. I let children enter the dataset at age 18, which is the most common age of graduation from high school in Denmark. As a result, one cannot discard the initial conditions problem in this context, since one does not know exactly which period is to be counted as the initial.

The data I apply for the extensive margin holds child-parent observations where the child studies actively for at least one year in higher education and resides in one of the major university cities[20], also for at least one year, while he is of age 18-30. I focus on children in education as these experience an anticipated and sustained fall income relative to their contemporaries. Conditioning on children who reveal interest in living in a city for at least one period is relevant since buy-to-let investments are overwhelmingly located within cities (where small, shareable apartments are mostly found).

Note that as shown in Appendix 6, the dynamic random effects probit model is based on a set of parameters accounting for the effects of i) the decision in $t-1$, ii) contemporary time-varying variables, iii) non-time varying variables used to model the random effect, iv) initial conditions and v) year dummies. For the following interpretation of parental behavior, only i)-iii) are of primary interest and so iv)-v) are relegated to tables in Appendix 6.

Results on extensive margin

Considering the qualitative effects on choice probabilities, a handful of intuitive results stand out in Table 4. First of all, there is strong state dependence for tenants in buy-to-let apartments, expressed by large and positive coefficient on lagged buy-to-let state. This outcome is to be expected due to the large fixed costs involved in trading real estate and moving address.

More interestingly, child income has a significant negative marginal effect on the buy-to-let probability. Note that since child income varies over time, it affects choice probabilities both through its contemporary realization and through the cross-sectional unit average, E_i[Child income], which enters into the parameterization of the random effect distribution.[21] Both coefficients are negative, implying that both permanent and temporary lower child income is

[20] Copenhagen, Aarhus, Odense, Aalborg, Esbjerg

[21] The random effect distribution also utilizes initial observations, see Appendix 6.

Table 4: Dynamic random effects probit of the buy-to-let investment decision for parents with children who receive higher education and live in a main city for at least one year between ages 18-30.

	Coefficients	S.E.
Constant	-3.4043***	0.0246
Buy-to-let$_{t-1}$	2.8508***	0.0091
Parental income$_t$	0.0411***	0.0020
Child income$_t$	-0.0514***	0.0028
$\mathbb{1}\{$Child in high educ$_t\}$	0.3853***	0.0064
$\mathbb{1}\{$Child age $\geq 24_t\}$	-0.1958***	0.0058
$\mathbb{1}\{$Self-employed parent$_t\}$	0.0272*	0.0139
$\mathbb{1}\{z^* <$ Parental income$_t\}$	-0.0221*	0.0114
Parental income$_t \times \mathbb{1}\{z^* <$ Parental income$_t\}$	-0.0082***	0.0022
$\mathbb{1}\{$Child has graduated$_t\}$	-0.1363***	0.0089
$\mathbb{1}\{$Parent is divorced$\}$	-0.3587***	0.0057
N properties owned	0.1539***	0.0023
N siblings	-0.1031***	0.0029
$\mathbb{1}\{$Parent w. high education$\}$	0.1233***	0.0054
$\mathbb{1}\{$Child does graduate$\}$	0.0181**	0.0073
E$_i$[Parental income]	0.0040**	0.0016
E$_i$[Child income]	-0.0237***	0.0043
E$_i[\mathbb{1}\{$Child in high educ$\}]$	-0.0992***	0.0151
E$_i[\mathbb{1}\{$Child age $\geq 24\}]$	0.1326***	0.0223
E$_i[\mathbb{1}\{$Self-employed parent$\}]$	0.1589***	0.0200
E$_i[\mathbb{1}\{z^* <$ Parental income$\}]$	0.2961***	0.0165
E$_i$[Parental income$\times \mathbb{1}\{z^* <$ Parental income$\}]$	0.0046*	0.0025
E$_i[\mathbb{1}\{$Child has graduated$\}]$	-0.1878***	0.0207
σ_a	0.5890	-
Log-likelihood	-250554.67	-
N	8662730	-
N_i	980572	-

*Notes: *** significant at 1% level, ** at 5%, * at 10%. Prices and incomes are in DKK 2015 and scaled by 100,000. Z-scores in parentheses. z^* denotes the cut-off between mid or bottom taxation (depending on the period). $E_i[v]$ denotes the average of variable v for cross-sectional unit i.*

associated with a higher chance of parents investing in a buy-to-let. Of course, one cannot make a causal statement here. It may be that parents are financially supporting children by means of housing, as in e.g. Kaplan (2012), or it may be that children work less to make ends meet when provided a cheaper housing solution. In either case, the important notion in terms of altruism is that parents are in fact providing buy-to-let apartments more often as child income falls.[22]

As explained in Section 2, the tax system incentivizes buy-to-let investments among high income parents because interest payments are deductible and because rent payments can be retained for taxation at a point later in life when income is lower. Controlling for parental income therefore includes an indicator for having income above the top tax cut-off, z^*, and an interaction term.

Unsurprisingly, higher parental income is associated with higher likelihood of investing and there is a pronounced additional effect of having income in the top tax bracket. Note that the positive effect of having income in the top tax bracket comes through the average realization entering the random effect, while the coefficient to the immediate effect is negative and much smaller.[23] The fact that average income is more important to the decision than the contemporary realizations is natural considering the long-term nature of a buy-to-let investment. We observe a negative coefficient to the contemporary interaction term between parental income and income above the top tax cut-off, but a positive coefficient to the average interaction. The interpretation is therefore not clear, but evidently, the interacted effect is also small.[24]

The overall picture is therefore that investment propensity increases with parental income

[22] For a more elaborate theory of family transfers involving moral hazard considerations, see Barczyk and Kredler (2014).

[23] In fact, since most units have at least 8 observations, the coefficient to $\mathrm{E}_i[\mathbb{1}\{z^* < \text{Parental income}\}]$ will swamp the effect of $\mathbb{1}\{z^* < \text{Parental income}_t\}$ even if there is only 1 period with income in top tax bracket.

[24] Note that a negative coefficient to the interaction would suggest a diminishing inclination to make investments with additional income above the top tax cut-off. This is in fact consistent with incentives laid out previously because additional income decreases the chance of ever getting below the top tax cut-off, making the option to retain rent obsolete. On the other hand, a positive coefficient suggests that the buy-to-let is a luxury good which is also a natural interpretation.

and decreases with child income. As children will normally pay a lower rent in a buy-to-let, such behavior falls well in line with parents being altruistically motivated and aiming to reduce welfare differences between themselves and their children.

In addition, there is a large positive effect of the child being enrolled in higher education. As discussed above, one interpretation of this effect is that parents are helping their children find housing on a market which has been particularly tight in urban areas close to higher education campuses. A related interpretation is that parents are helping children smooth consumption intertemporally, since they may accommodate fluctuations in child income when setting the rent. For students this may involve setting a low rent while enrolled in studies and then increasing it upon graduation.[25]

There is a marked association of buy-to-let investments with the number of properties owned by parents. This is likely picking up both a wealth effect and the level of financial sophistication of parents. Being divorced is associated with lower chance of buy-to-let investments, which can be rationalized with the notion that less marital stability also offers less personal resources for time consuming investments.

The year-dummies reported in Table 9 in Appendix 6 show that buy-to-let investments have closely followed the general housing cycle with peaking activity in the mid 2000s, a steep decline after the financial crisis and a subsequent recovery. This suggests a level of momentum trading among parents, but it may also reflect a wealth effect stemming from similar swings in their housing equity. Adding additional years to the panel will likely be informative of this question, as Copenhagen housing prices have surged in the period 2013-2018 but remained stagnant elsewhere (as opposed to nation wide boom in the 2000s).

Finally, a set of regressions on alternate subpopulations are presented in Appendix 6 which do not lead to different conclusions.

[25]Future editions of the paper will dig into this perspective.

Table 5: Linear regressions of buy-to-let purchase price on family characteristics.

	ln(Price)		
	Copenhagen	Aarhus	Elsewhere
Constant	1.4752***	1.6703***	1.3503***
	(59.00)	(60.45)	(44.36)
Child income	-0.0657***	0.0243	0.0288**
	(-6.17)	(1.41)	(2.07)
Child income^2	0.0197***	-0.0011	0.0019
	(7.18)	(-0.23)	(0.57)
Parental income	0.0113***	0.0270***	0.0382***
	(3.75)	(5.92)	(7.22)
Parental income^2	0.0006***	-0.0008***	-0.0009***
	(3.30)	(-3.04)	(-2.73)
Parental bank deposits	0.0034***	0.0021***	0.0037***
	(9.93)	(5.10)	(5.95)
$\mathbb{1}${Firm tax}	0.0389***	0.0655***	0.0412***
	(6.22)	(6.19)	(3.98)
$\mathbb{1}${Capital returns tax}	-0.0111	0.0091	-0.0125
	(-1.53)	(0.77)	(-1.13)
$\mathbb{1}${Parent w. high education}	0.0676***	0.0562***	0.0612***
	(11.87)	(6.42)	(7.03)
$\mathbb{1}${Parent is divorced}	0.0167***	-0.0007	-0.0144
	(2.85)	(-0.07)	(-1.39)
N properties owned	0.0211***	0.0134***	0.0017
	(8.51)	(4.09)	(0.41)
$\mathbb{1}${Child age $\geq$ 24}	0.0523***	0.0332***	0.0266**
	(8.64)	(3.55)	(2.51)
$\mathbb{1}${Child in short education}	-0.0194*	-0.0206	-0.0380***
	(-1.86)	(-1.21)	(-2.84)
$\mathbb{1}${Child in high educ}	0.0409***	0.0313***	0.0315***
	(7.07)	(3.65)	(3.07)
$\mathbb{1}${Child has graduated}	0.0010	-0.1408***	0.0366
	(0.01)	(-2.59)	(0.57)
$\mathbb{1}${Child graduates before 30}	0.0372	0.1325***	-0.0437
	(0.58)	(2.70)	(-0.70)
R^2	0.60	0.50	0.36
N	18597	6710	8548

*Notes: *** significant at 1% level, ** at 5%, * at 10%. Prices and incomes are in DKK-2015 and scaled by 100,000. Robust Z-scores in parentheses. Year-dummies are omitted, but found in Appendix 6, Table 12. Exogenous variables correspond to the year of buy-to-let purchase. Data includes only buy-to-let cases where children are enrolled in higher education between ages 18-30. Bank deposits reflect all liquid savings held by December 31st each year.*

4.2 Intensive margin of buy-to-let investments

The empirical behavior on the intensive margin of the investment choice is described a linear regression of apartment market value on family characteristics, which is presented in Table 5.[26] The results confirm the intuitions for altruistic behavior and incentives discussed in Section 2.

As depicted in Figure 2, housing prices have followed different trajectories in Copenhagen, Aarhus and the remaining Denmark. The regression is therefore run separately on buy-to-let observations in each region and the investment behavior appears quite stable across regions. For interpretability, I include only cases where the child is enrolled in further education between ages 18-30.

As expected, parental income has a positive effect on the decision of the buy-to-let value. More surprisingly, the value of chosen buy-to-let apartment is practically linear in parental income, since coefficients to the squared income terms are negligible. Recent papers by Lockwood (2018) and Barczyk et al. (2019) show that gifts and bequests are best thought of as luxury goods. As such, one may have expected a non-linear increase in apartment value with income. The fact that this is not observed might be attributed to the increasing rental payments that the child must bear with increasing apartment value. As discussed previously, parents are only able to cut rent to a certain extent, meaning that they will not be helping the child by providing an overly luxurious apartment.

Contrary to the extensive margin, the effect of child income is predominantly positive. The effect of contemporary income is U-shaped in Copenhagen and practically linear in Aarhus and elsewhere. More importantly, there are clear effects of the age of children their education choice as both of these are related to future income. Higher age means that the child is closer to entering the labor market full-time while higher education increases future income (and the reverse is true for short education). The motive of helping children to smooth consumption may be read into this, especially because of the prospect of children

[26] Since parents overwhelmingly only invest in a buy-to-let once, panel data techniques are not available.

buying the apartment from parents after graduation. Also, parents and children may agree to let the rent increase over time within the bounds of rent regulation.

As explained in section 2.1, parents can benefit from subscribing to one of the tax schemes for self-employment, in part by avoiding taxation of rent payments by through income shifting and in part through the deductibility of interest expenses. Both of these options work to subsidize the investment and, consequently, parents using the firm taxation scheme also opt for more expensive investments. Recall that the the capital returns scheme was easier to manage, but also less lucrative in terms of interest deductibility and income shifting. It is therefore not surprising that this scheme does not seem to entail more expensive investments.

Finally, parental wealth is arguably a key determinant of the buy-to-let investment due to the down payment requirement. Only a crude measure of parental wealth is unfortunately available for a major part of the data horizon, namely properties owned and end of year bank deposits, but these do exhibit the expected positive association with apartment value.

4.3 Reselling the buy-to-let

Lastly we consider the behavior involved in reselling a buy-to-let apartment when it is no longer rented to a child. As explained in Section 2, parents may keep the apartment and rent it out to a third party, but due to regulation, they are more likely to sell it - possibly to the child a discount. Table 13 in Appendix 6 shows the number of buy-to-let apartments sold by parents annually, conditioning on selling to child or third party. There is a clear pro-cyclical movement in both total sales and the share that goes to children, indicating that both investment motives and gift motives may be at play.

To shed light on the presence of altruistic motives when reselling buy-to-let apartments, I use a measure of an apartment's development in trading price relative to its market value, i.e. its markup, as it is handed off by the owning parent. This measure is regressed against indicators for selling to children, family income and housing cycle conditions. One needs the measure of relative price development because we do not observe the counterfactual

Table 6: Linear regression of relative resale price of buy-to-let appartments

	$\tilde{\Delta}$Price
Constant	-0.0666***
	(-4.24)
Parental income$_t$	-0.0053***
	(-3.54)
Child income$_t$	-0.0095***
	(-2.68)
$\mathbb{1}$\{Selling to child\}	-0.2385***
	(-11.06)
Parental income$_t$ $\times$ $\mathbb{1}$\{Selling to child\}	-0.0021
	(-0.84)
Child income$_t$ $\times$ $\mathbb{1}$\{Selling to child\}	0.0175***
	(2.71)
$\mathbb{1}$\{Selling during slump\}	-0.0543***
	(-4.35)
$\mathbb{1}$\{Buying during slump\}	0.0327**
	(2.41)
$\mathbb{1}$\{Selling to child\} $\times$ $\mathbb{1}$\{Selling during slump\}	0.1642***
	(7.00)
$\mathbb{1}$\{Selling during boom\}	-0.2519***
	(-27.10)
$\mathbb{1}$\{Buying during boom\}	0.0909***
	(10.20)
$\mathbb{1}$\{Selling to child\} $\times$ $\mathbb{1}$\{Selling during boom\}	-0.0152
	(-0.88)
$\mathbb{1}$\{Parent is divorced\}	0.0393***
	(3.61)
$\mathbb{1}$\{Parent is single\}	-0.0171
	(-1.47)
N properties owned	0.0114***
	(4.39)
$\mathbb{1}$\{Parent w. high education\}	-0.0405***
	(-5.48)
BTL size, m2	-0.0002
	(-1.22)
$\mathbb{1}$\{BTL in CPH\}	-0.0228***
	(-3.14)
R^2	0.20
N	11737

*Notes: *** significant at 1% level, ** at 5%, * at 10%. Incomes are in DKK-2015 and scaled by 100,000. Robust Z-scores in parentheses.*

free-market price of an apartment that is sold within the family. We do however need the comparison when identifying rebates to children. I therefore construct a measure of an apartment's relative change in price, comparing to local trends, between the time of initial purchase by parents and the time of reselling.

Letting p_s denote the price an apartment is observed to sell at in period s, let p_b denote the price an apartment is purchased at in period b and let $\bar{p}_z$ be the median apartment price pr m^2 of zone[27] z in which the the buy-to-let apartment is situated. The relative change in trading price of an apartment is then

$$\tilde{\Delta}\text{Price} = \frac{p_s}{p_b} - \frac{\bar{p}_{z,s}}{\bar{p}_{z,b}}, \tag{1}$$

which is the LHS of the regression in table 5. A $\tilde{\Delta}\text{Price} > 0$ implies that the apartment increased more in value than the median of m^2 prices in its close vicinity to the benefit of the seller.

The regression demonstrates that children are indeed offered a significant rebate when buying from their parents, consistent with a gift motive of parents. The coefficient on selling-to-child shows that parents all else equal let go of 24 percentage points of the local price development when selling to their children. This does represent a sizeable gift for those who sold at the peak of the boom (or in recent years), as prices increased by almost 300% between 1995-2006 in Copenhagen.

Interestingly, child income has a positive effect when interacted with selling-to-child, so that higher income children are given a smaller rebate on the apartment. As will be demonstrated in the model section, altruism of parents imply that they will try to equalize welfare within the family. Cutting the rebate for higher income children is therefore clearly consistent with a model of parental altruism. Followingly, one may also expect the interaction

[27] I observe the commuting zone each housing sale in the sales register. Commuting zones are fine grained in urban areas with a median size of 2 km^2. The growth rate of median prices should therefore be a good approximation of apartment's expected change in price in absence of any idiosyncratic improvements or impairments.

term with parental income to be negative. It is however insignificant, but this may be caused by the amount of variation in data and that the selling-to-child indicator runs in the same direction.

The regression also controls for macro conditions during the time of selling and buying the apartments. The Danish housing market experienced a boom period in 2003-2007 and a slump in 2008-2012. Interacting the indicator for selling during slump with selling to child reveals that parents largely undid the rebate when prices were declining. It seems therefore that parents were primarily sharing *gains* on the apartment with children rather than providing an unconditional gift.

5 A collective model of parental buy-to-let

The aim of the model is to describe the buy-to-let decision process from the view point of a partly altruistic parent. For simplicity, there are two periods and perfect foresight with respect to incomes and prices. The problem of the parent has a discrete-continuous structure: when investing in a buy-to-let apartment, she also decides on its size and the down payment. In the beginning of period 2, the child then buys the apartment from the parent at a rebate equal to the down-payment. When the parent does not invest in a buy-to-let, she merely decides on her own consumption. Importantly, the child has the option to reject the buy-to-let apartment offered by the parent and go find a rental on his own. Whatever offer the parent comes up with must therefore yield higher utility for the child than his outside option. Hence, the structure of the model is as follows: the parent first realizes how well off the child, and herself respectively, would be without investing in a buy-to-let. Then she calculates her utility of making the optimal investment conditional on the investment being an improvement for the child and finally makes the decision by comparing the two potential outcomes.

Following the terminology of the family economics literature, see e.g. Browning et al. (2014), the parent has *caring* preferences for the child, which means that the child's outcome utility factors into the utility perceived by the parent through the coefficient $\mu \in [0, 1]$. A standard result, also discussed in Browning et al. (2014), is that the outcome must, as a consequence of caring preferences, be Pareto efficient for the family members. Intuitively, μ is proportional to the tangent of the utility possibility frontier at the solution of the problem. As explained, a buy-to-let can both serve as a way of keeping income within the family, a tax management instrument, an easier access to housing, a way of transferring wealth but is also costly for the parent (since she is giving up her private wealth). The size of μ thus determines how much these trade-offs are tilted in favor of the child.

5.1 Separate decisions

Given the process of backward induction in the solution of the model, we will begin with the separate problems of child and parent if no buy-to-let apartment is purchased. I assumed that parents and children are *not* interacting or making any transfers to one another, and the decisions of the parent is thus not directed by altruism towards the child. The parent does not consider moving so her only objective is to plan consumption subject to income and initial wealth. The objective of the child is to plan both consumption and housing.

Parents

The parent receives a stream of exogenous taxable labor income in each period, denoted z_{pt}. In line with Danish tax law after 2008, income is subject to a tax system with two brackets. The marginal tax in the lower bracket is denoted τ_b and the marginal tax in the upper bracket is given by $\tau_b + \tau_u$. The two brackets are separated by the point z^*. After-tax income $\Gamma(z)$ is thereby

$$\Gamma(z) = z - (\tau_b z + \tau_u \max(z - z^*, 0)). \tag{2}$$

The parent owns from the outset a freehold of size h_p, measured in m^2, which she has borrowed against to finance the purchase. In the model, freeholds trade at a constant price pr m^2, p^m, so that the market value of a freehold $P(h)$ is

$$P(h) = p^m \times h. \tag{3}$$

For simplicity, I assume that the parent has financed the home purchase using an interest only mortgage with a loan-to-value ratio of λ_0.[28] The initial wealth of the parent therefore amounts to $(1 - \lambda_0)P(h_p)$ in home equity. I assume that they can costlessly readjust their leverage by cashing out or paying off debt, corresponding to a change in λ.

The Danish mortgage system dictates that only 80 pct of real estate's market value can serve as collateral for mortgage bonds. The remaining need for financing is covered by a bank loan, normally with a markedly higher interest rate. Let λ^* denote the 80 pct limit, r^m the mortgage bond rate and r^s the spread between the bank loan rate and the mortgage bond rate. The period-wise mortgage expense $\mathcal{R}$ therefore depends on the loan-to-value ratio and the purchase value of the home

$$\mathcal{R}(\lambda, P(h)) = P(h)(r^m\lambda + r^s \max(\lambda - \lambda^*, 0)). \tag{4}$$

When solving the model numerically, the max operator is replaced with the softmax for better convergence.[29] Note that the softmax operator may alternatively be interpreted as a measure of the bank's uncertainty of the exact current market value of the home (and hence the loan-to-value).

Denoting the consumption of parents c_{pt}, the intertemporal budget constraint of a parent

[28]IO mortgages are not exactly the representative kind of mortgage for parents engaging in buy-to-let investments, but I focus on this type to avoid muddling the analysis unnecessarily by installment payments.

[29]Mortgage costs when softmax is implemented are given by

$$\mathcal{R}(\lambda, P(h)) = P(h) \int_0^\lambda r^m \Big(1 - \frac{e^{(\tilde{\lambda}-\lambda^*)/\sigma^s}}{1 + e^{(\tilde{\lambda}-\lambda^*)/\sigma^s}}\Big) + (r^m + r^s)\frac{e^{(\tilde{\lambda}-\lambda^*)/\sigma^s}}{1 + e^{(\tilde{\lambda}-\lambda^*)/\sigma^s}} d\tilde{\lambda} \tag{5}$$

The scaling parameter σ^s is set sufficiently low so that the smoothened interval around λ^* effectively spans less than 2 pct points on each side.

is given by

$$c_{pt} + \mathcal{R}(\lambda_t, P(h_p)) = \Gamma(z_{pt}) + (\lambda_t - \lambda_{t-1})P(h_p) \tag{6}$$

Since the model essentially is a partial life-cycle model, the parent and child must take into consideration the value of their remaining assets b_{jt} at the end of period 2. Remaining assets, or terminal wealth, is for parents given by home equity at the end of period 2.

Parent and child are assumed to have similar preferences for housing and consumption which is denoted $u_j(c_{jt}h_{jt})$ and subscripted by $j = p, k$ for parents and children. Similarly, they share preferences for terminal wealth given by the function $v(b_j)$. Note that although parents do not actually make adjustments to their own housing position, h_p still enters into their utility. We need this feature since the parent's caring preferences would not be meaningful if the utility functions of family members have different structures.

We are now in the position to state the problem faced by parents in the situation where no buy-to-let investment is made. Let W^p be associated with the indirect utility of a parent making no buy-to-let investment, $\boldsymbol{c_p} = \{c_{p1}, c_{p2}\}$ and $\boldsymbol{\lambda_p} = \{\lambda_{p1}, \lambda_{p2}\}$, the problem is

$$W^p = \max_{\boldsymbol{c_p}, \boldsymbol{\lambda_p}} \{u_p(c_{p1}, h_{p1}) + \beta u_p(c_{p2}, h_{p1}) + \beta v_p(b_{p2})\} \tag{7}$$

$$s.t.$$

$$c_{p1} + \mathcal{R}(\lambda_1, P(h_p)) = \Gamma(z_{p1}) + (\lambda_1 - \lambda_0)P(h_p) \tag{8}$$

$$c_{p2} + \mathcal{R}(\lambda_2, P(h_p)) = \Gamma(z_{p2}) + (\lambda_2 - \lambda_1)P(h_p) \tag{9}$$

$$b_{p2} = p_h h_p(1 - \lambda_{p2}) \tag{10}$$

$$\lambda_{p1}, \lambda_{p2} \in [0, 1) \tag{11}$$

$$\lambda_{p0}, h_p : \ given \tag{12}$$

Children

When no buy-to-let investment is made, children must find a rental on their own. I assume they can do so in each period without incurring search costs[30] or advance payments. To make utility comparable across the no-investment and investment scenario, I let the child make savings b_k for terminal wealth in both periods. As will be explained below, the parent transfers an amount of terminal wealth to the child through the buy-to-let in the investment scenario. Letting the child make an active saving decision in the no-investment scenario therefore increases compatibility between the two.

The market for rentals is assumed to display a marginally decreasing rent pr m^2 as apartments grow larger. An illustration of the rent function is given in Figure 8 in Appendix 6. The notion that marginal rent pr. m^2 is decreasing with the size of a rental is intuitive since smaller apartments imply more people in the same building, leading to higher maintenance and service costs pr. building for a landlord. This relationship seems undescribed in the academic literature, but it is clearly reported for Danish rentals in Hansen and Iversen (2017). Therefore, I will assume that rent includes a service cost, $K(h+\bar{h})^{-\eta}$, which is a decreasing function of h above a minimum size, $\bar{h}$. It also includes a fraction χ of the home's market value as a freehold, $P(h)$. The annual rent r for an apartment of size h is therefore

$$r(h, P(h)) = K(h+\bar{h})^{-\eta} + \chi P(h). \tag{13}$$

Note that this specification of rent has the advantage that I can let the parent forgo the service cost in the investment scenario, creating a way for the parent to make an indirect transfer to the child.

Finally, the child receives an exogenous labor income z_{kt} which is liable to taxation as in equation (2). Letting bold symbols denote the vectors of period wise choices, the problem the child comes down to

[30]Including a search cost is an obvious expansion of the model for future work.

$$W^k = \max_{\boldsymbol{c_k},\boldsymbol{h_k},\boldsymbol{b_k}} \{u_k(c_{k1}, h_{k1}) + \beta u_k(c_{k2}, h_{k2}) + \beta v_k(b_{k1} + b_{k2})\} \tag{14}$$

$$s.t.$$

$$c_{k1} + r(h_{k1}, P(h_{k1})) + b_{k1} = \Gamma(z_{k1}) \tag{15}$$

$$c_{k2} + r(h_{k2}, P(h_{k1})) + b_{k2} = \Gamma(z_{k2}) \tag{16}$$

5.2 Buy-to-let investment

We now turn to the problem of a parent choosing to invest in a buy-to-let. The structure of the problem changes as the decision of the parent now influences the child in the following way: at the beginning of period 1, the parent uses a share of her endowed home equity to make a down payment of size γ on an apartment which she then lets to the child for the rest of period 1. The apartment is of size h_k. In charging rent, she renounces the service fee usually associated with renting in equation (13) as a matter of parental altruism. Simultaneously, she decides how large a fraction of the rental income is to be extracted for consumption in period 1 and how much is retained for period 2.

In the beginning of period 2, the parent sells the apartment to the child at a discount equal to the down payment. Although the coincidence between the discount and the down payment is perhaps arbitrary, it is a handy way of reducing the complexity of the model without changing the motives qualitatively. Recall from the discussion of the tax system in Section 2 that large gifts are liable to taxation but in offering a discount on the apartment, the parent can make an indirect untaxed wealth transfer of $\gamma P(h_k)$ to the child. The transferred wealth in form of home equity is now kept by the child through period 2 and appears as terminal wealth. Since the child purchases the home at the beginning of period 2, it must pay the mortgage costs of a loan of size $(1 - \gamma)P(h_k)$ in this period.

The full problem is stated in equations (17) - (26) below, which I will go through successively. The welfare of the child factors into the utility function of a parent through the caring parameter μ, which gives the objective function in (17). The optimal plan of the parent consists in choosing her level of consumption in each period as before, the size of the buy-to-let apartment, h_k, how large a fraction of rent income to retain between period 1 and 2, w, and the loan-to-value on the buy-to-let apartment, γ. Since the child does not save up between period 1 and 2, his consumption is determined through the rental and mortgage costs, which are a function of apartment size and down payment.

$$V = \max_{\boldsymbol{c_p}, h_k, w, \gamma} \{u_p(c_{p1}, h_p) + \beta u_p(c_{p2}, h_p) + \beta v_p(b_p) \\ + \mu[u_k(c_{k1}, h_k) + \beta u_k(c_{k2}, h_k) + \beta v_k(b_k)]\} \quad (17)$$

$s.t.$

Constraints on parent's budget

$$s = P(h_k)(\chi - \psi) - \mathcal{R}(1-\gamma, P(h_k)) \quad (18)$$

$$c_{p1} + \mathcal{R}(\lambda_{p1}, P(h_p)) = \Gamma(z_{p1} + ws) \quad (19)$$

$$c_{p2} + \mathcal{R}(\lambda_{p1}, P(h_p)) = \Gamma(z_{p2} + (1-w)s) \quad (20)$$

$$(\lambda_{p1} - \lambda_{p0})P(h_p) = \gamma P(h_k) \quad (21)$$

$$b_p = (1 - \lambda_{p1})P(h_p) \quad (22)$$

Constraints on child's budget

$$c_{k1} + \chi P(h_k) = \Gamma(y_{k1}) \quad (23)$$

$$c_{k2} + \mathcal{R}(1-\gamma, P(h_k)) = \Gamma(y_{k2}) \quad (24)$$

$$b_k = \gamma P(h_k) \quad (25)$$

$$W^k \leq u_k(c_{k1}, h_k) + \beta u_k(c_{k2}, h_k) + \beta v(b_k) \quad (26)$$

Constraints on financial variables

$$w \in [0,1], \ \lambda_{pt} \in [0,1), \ \gamma \in (0,1]$$

Let s in equation (18) denote the net rent income that accrues to the buy-to-let enterprise in period 1. It consists of the fraction χ of the apartment value, subtracted the maintenance costs ψ and mortgage costs $\mathcal{R}(1-\gamma, P(h_k))$. Recall that these were fully deductible under the firm taxation system. Note that maintenance and mortgage costs are important to the problem structure as these are money which flow out of the hands of the family irrespective of making the buy-to-let investment or not. The incentive to make the investment therefore falls with their relative magnitude.

In order to optimize intertemporal consumption and tax payments, a parent chooses to withdraw a share w of s in period 1 and retains $1-w$ for period 2. This, together with exogenous labor income, constitutes the RHS of parent's budget constraints (19)-(20).

As noted above, the parent uses part of her home equity to make the down payment on the buy-to-let apartment at the start of period 1. Since the down payment equals the discount given to the child, there is no wealth gain for the parent associated with selling. Equation (21) expresses exactly that the leverage of parents increases by the down payment made on the buy-to-let in period 1 and stays constant thereafter. Parents must therefore pay the costs of their existing mortgage and the down payment in addition in both periods. This explains the LHS of (19)-(20). Finally, the parent is left with $1-\lambda_{p1}$ of her home value in terminal wealth by equation (22). The fact that the parent makes a trade-off between how much terminal wealth she will end up with, and how much is to be passed on to the child through the down payment, underlines the importance of the wealth motive $v(b)$.

The budget constraints of the child, (23) - (24), consists of his after-tax income on the RHS. The LHS includes rent payments (period 1) and mortgage costs (period 2) since the apartment is purchased by the child at the beginning of period 2. Equation (25) states that the child is left with the down payment as terminal wealth. Lastly, the inequality (26) captures that the child has the option of rejecting the buy-to-let offered by the parent if his implied utility falls below the value of finding a rental by himself, W^k.

Discrete choice of making a buy-to-let investment

The final discrete-continuous choice of making a buy-to-let investment or not comes down to comparing the value of making the investment, V, and the value of relying on separate decisions. Because the utility of children enters additively into the utility function of parents through the caring parameter, the value of running separate decisions is simply $W^p + \mu W^k$ from the perspective of the parent. Denoting the policy with respect to the discrete choice of making an investment by $\mathbb{I}^* = \{\text{Invest, not invest}\}$, this becomes

$$\mathbb{I}^* = \arg\max(V, W^p + \mu W^k). \tag{27}$$

For the exposition (and eventual estimation) it is typically more convenient to express the policy probabilistically. This is here obtained by assuming that each choice involves a random utility component $\varepsilon_i \sim EV(0, \sigma_\varepsilon)$ that expresses benefits whose realization are not known to the parent at the time of deciding. As the parent now must integrate out the utility shocks when choosing to invest or not[31], her utility maximizing policy becomes a randomization strategy with the propensity to choose the buy-to-let investment

$$\text{Prb}(\mathbb{I}^* = \text{Invest}) = \frac{\exp(V/\sigma_\varepsilon)}{\exp(V/\sigma_\varepsilon) + \exp((W^p + \mu W^k)/\sigma_\varepsilon)}. \tag{28}$$

Preferences

The parent and child have similar preferences for housing, consumption and terminal wealth. Housing and consumption is aggregated by the Cobb-Douglas function as per the commonly observed constant budget share of housing, see e.g. Davis and Ortalo-Magné (2011), and embedded in the CRRA utility function

$$u_j(c_{jt}, h_{jt}) = \frac{\left(c_{jt}^{1-\phi} h_{jt}^{\phi}\right)^{1-\sigma}}{1-\sigma}. \tag{29}$$

[31] See eg. Train (2009) for a derivation.

The utility of terminal wealth is likewise given by the CRRA function and scaled by $\iota > 0$, so that

$$v_j(b_j) = \iota \frac{(b_j)^{1-\sigma}}{1-\sigma}. \tag{30}$$

Discussion of the model

In taking stock of the model at its current state, it is worth mentioning where future development would be warranted for a model to fit data closely. First of all, a longer horizon than two periods is critical to expose the real tax-benefit of retaining rent earnings, simply because it increases the total amount and because parents then have a larger chance experiencing a fall in future labor income. Second, stochasticity in future prices and income are important to a home investment, especially considering the large fluctuations seen on the Copenhagen housing market. Third, there is no disutility associated with managing the buy-to-let for parents and neither is there any fixed cost when trading with the apartment. Adding these to the problem would also be necessary to obtain a proper fit.

5.3 Calibration

I calibrate the model so that it applies to conditions in Copenhagen 2012 in order to evaluate it on a data sample for that year. The joint distribution of state variables in the sample is an approximation to the data on which the discrete choice regressions were applied, see Appendix B for details. Fitting the model to data will be a topic for future work, but for now I fix parameters externally so as to illustrate predictions from the model.

All parameters can be found in Table 7. I set the preference for housing, ϕ, to match the average expenditure share on housing among renters in Copenhagen. Hansen and Iversen (2017) report expenditure shares concentrated around 0.25 with the 75th percentile reaching 0.3. Based on Hjelmar (2009), young adults spend a comparatively large share of income on housing, so I let $\phi = 0.3$. In line with related papers such as Sommer and Sullivan (2018) or Barczyk et al. (2019), the relative risk aversion parameter, σ is set to 2. The altruism

Table 7: Parameter settings

Parameter		**Value**
Preferences		
Weight on children utility	μ	0.4
Preference for housing	ϕ	0.3
Risk aversion, total consumption	σ	2.5
Weight on terminal wealth	ι	0.01
Tax system		
Kink point between low and high tax bracket	y^*	400
Base tax	τ^b	0.4
Top tax	τ^u	0.15
Housing and mortgage costs		
Marginal price of housing	p^m	25
Capital costs of renting	χ	0.05
Maintenence costs of renting	ψ	0.01
Service cost of renting, coefficient	K	150
Service cost of renting, exponent	η	0.8
Service cost of renting, constant	$\bar{h}$	10
Mortgage backed interest rate	r^m	0.03
Spread between bank rate and mortgage rate	r^s	0.03
LTV limit on mortgage loan	λ^*	0.80
Softmax scaling parameter for mortgage rate	σ^s	0.02
Scale on utility shocks	σ^ε	0.01

Notes: Income and prices (incl. y^* and p^m) are measured in DKK1000s.

parameter μ varies substantially across studies that have formed an estimate of it; Barczyk et al. (2019) arrives at 0.4 while Kaplan (2012) at 0.04. Setting $\mu = 0.1$ is however enough to yield proper variation in investment probabilities when evaluating the model on data. The weight parameter on terminal wealth ι amplifies the child's value of getting a buy-to-let through the associated home equity. When $\iota = 0.04$, it ensures positive transfers between parents and children for the given value of μ.

The parameters of the tax system are readily available on the website of the Danish Ministry of Taxation[32] Likewise, the interest rate on mortgages and the spread to bank lending,r^m, r^s, are based on the average effective rates for 2012 reported by the central bank of Denmark.[33] The LTV limit $\lambda^* = 0.80$ is defined by Danish mortgage regulation.[34]

I choose the parameters for the service costs in the rental function such that rent pr. m^2 exhibits the shape found in Hansen and Iversen (2017). At the same time, I set $\chi = 0.05$ so that average rent pr. m^2 is on par with the average rent for unsubsidized student housing.[35].

5.4 Properties of the collective model

Based on the calibration, we are now in a position to consider the key properties of the model and how they relate to the empirical findings. Given the current state of the calibration, only qualitative comparisons are warranted. On this level, however, the implications of the model is mainly on par with the empirics.

Figure 4 shows the behavior of the model when varying either parental or child income while keeping everything else fixed. Panel (a) and (b) show that investment probabilities are increasing in parental income and decreasing in child income. This was indeed observed from the dynamic discrete choice regression in Table 4. Recall that the opposite effects are at

[32]Rates: www.skm.dk/skattetal/satser/tidsserier/centrale-skattesatser-i-skattelovgivningen-2010-2017
Cut-offs: www.skm.dk/skattetal/satser/tidsserier/centrale-beloebsgraenser-i-skattelovgivningen-2010-2017

[33]Danmarks Nationalbank: Registers DNRIURQ and DNRUGPI.

[34]The mortgage regulation is published at www.retsinformation.dk/eli/lta/2018/1188

[35]Reported in Municipality of Copenhagen (2019). Note that the return on rental properties have varied between 3.5%-5% for housing investments in Copenhagen, see e.g. the market analysis by Colliers (2019), which is consistent with these parameters.

Figure 4: Buy-to-let decision policies given child and parental income.

(a) Parent income and probability of buy-to-let

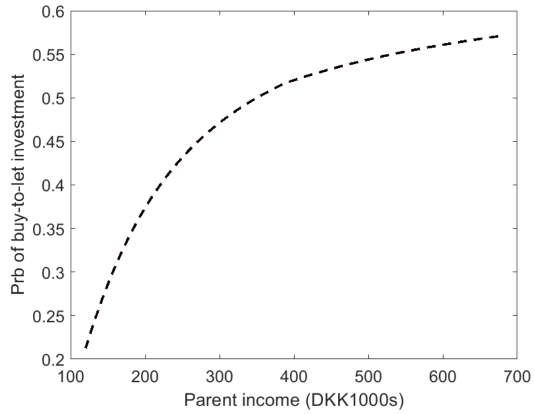

(b) Child income and probability of buy-to-let

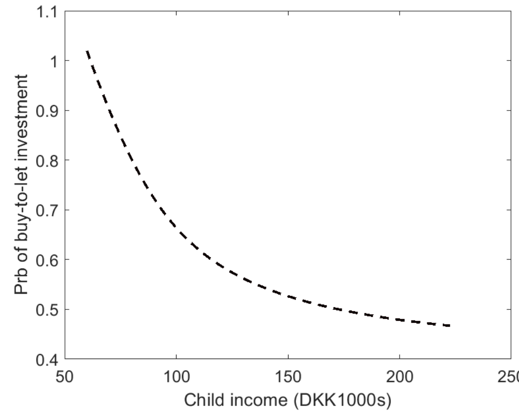

(c) Child income and size of of buy-to-let

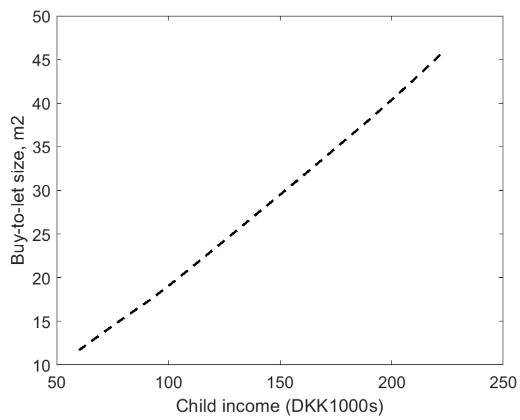

(d) Relative incomes and down payment-transfer

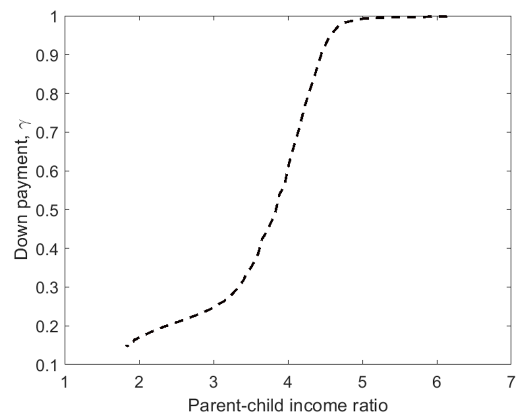

(e) Parent income and size of buy-to-let

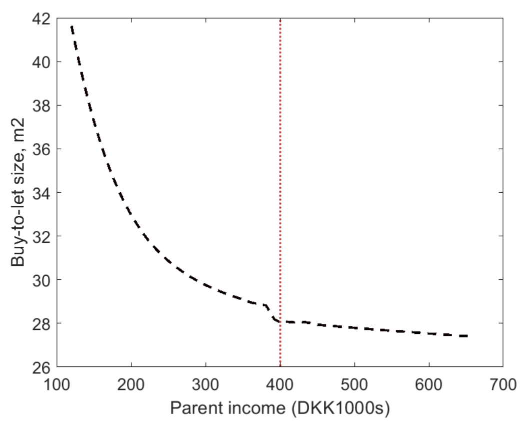

(f) Retained rent in period 1

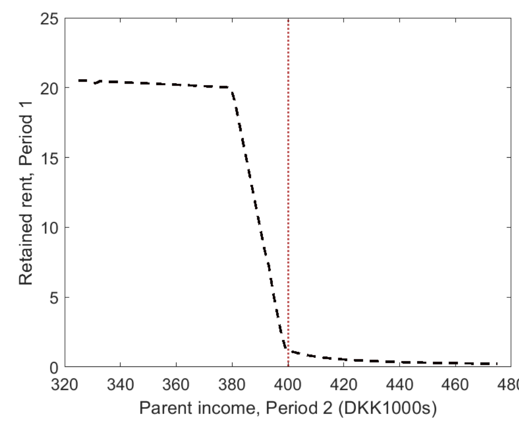

Notes: all variables, except the variables on the x-axes, are fixed at their average values in the approximated data sample when evaluating the model. That is, each diagram displays a change in the variable on the x-axis only. In figures (a)-(e), income is constant for parent and child across periods. In (d), child income is varied while parent income is fixed. In (f), parent income is set to 400 (the tax kink) in period 1. Both period 1 and period 2 income (x-axis) is the exogenous labor market income and does not include rent.

the heart of the collective model. A parent can use the buy-to-let to make a transfer to the child, which renders it a luxury good in so far $\mu < 1$ and marginal utility of consumption is decreasing. Since a buy-to-let decreases child housing costs and comes with the option to provide terminal wealth, lowering child income will increase the propensity of the parent to make a buy-to-let investment.

The demand for square meters in the buy-to-let, and thus rent, is linearly increasing in child income in panel (c). This is mainly a product of Cobb-Douglas preferences of housing and consumption. Panel (d) shows the more interesting relationship between the relative incomes of parent and child and the chosen level of wealth transfer through the down payment. Since parent income is fixed here, increasing the parent-child income ratio implies a poorer child. The CRRA preference for terminal wealth, equation (30), implies that the marginal value of terminal wealth increases exponentially with falling child income since any savings must be financed out of child income in absence of a parental wealth transfer. The value of making such a transfer through the rebate on the buy-to-let therefore increases correspondingly for the parent giving rise the non-linear relationship between parent-child income ratios and rebates/down payments.

Panel (e) in Figure 4 displays the chosen size of the buy-to-let when varying parent income while holding child income fixed (at 148). It is important to note that the figure shows the policy of the parent when conditioning on making an investment and it is likely that a parent would not find it optimal to invest at a large fraction of the incomes displayed. Nevertheless, the relationship between parental income and buy-to-let m^2 is uniformly negative. The reason for this is that the buy-to-let is a source of net income for the parent that increases with apartment size. A parent with low income and much wealth will therefore ask the child to accept a larger buy-to-let so as to generate more income for her.[36] This mechanism is not found in data, but it may not necessarily invalidate the model. For one, we do not have observations of parents with large amounts of (illiquid) wealth and little income. Secondly,

[36]Note how the progressive tax system discourages this tendency in that there is a minor cliff in m^2 just at the cut-off to the top tax bracket.

the current model omits fees, fixed costs and running expenses associated with the buy-to-let investment, which would normally be financed out of current parental income. Including these may change the picture.

Finally, the panel (f) shows how bunching around the top tax cut-off will occur in the model through retained earnings of the parent. In the diagram, a parent receives labor income in period 1 such that she is exactly at the top tax cut-off of 400. The diagram then shows how the decision to retain earnings depends on period 2 labor income. When period 2 labor income is lower than period 1 income by a large margin, the parents chooses to retain everything in period 1, as it is cheaper to get it taxed in period 2. From the point where rental income equals the distance between period 2 income and the top tax cut-off, the parent reduces retained earnings one for one with increased labor income. Since labor income is distributed smoothly around the top tax cut-off, the implication of such behavior will be a bunching around the top tax cut-off in the combined measure of rental and labor income of parents. Recall that this was indeed observed among owners of buy-to-let apartments in Figure 1.

In order to display how welfare and housing outcomes change with the introduction of buy-to-let investments, I evaluate the calibrated model on the sample of families described in Appendix 6. The densities of indirect utility of parents who, respectively, choose to or refrain from investing are displayed in Panel (a) of Figure 5. The density of indirect utility for parents who do invest is more concentrated and has a slimmer left tail than the density for non-investors, indicating that they are generally better off.

Figure 5: Distributions of welfare and apartment size in baseline scenario

(a) Density of welfare. (b) Densities of apartment size, m^2.

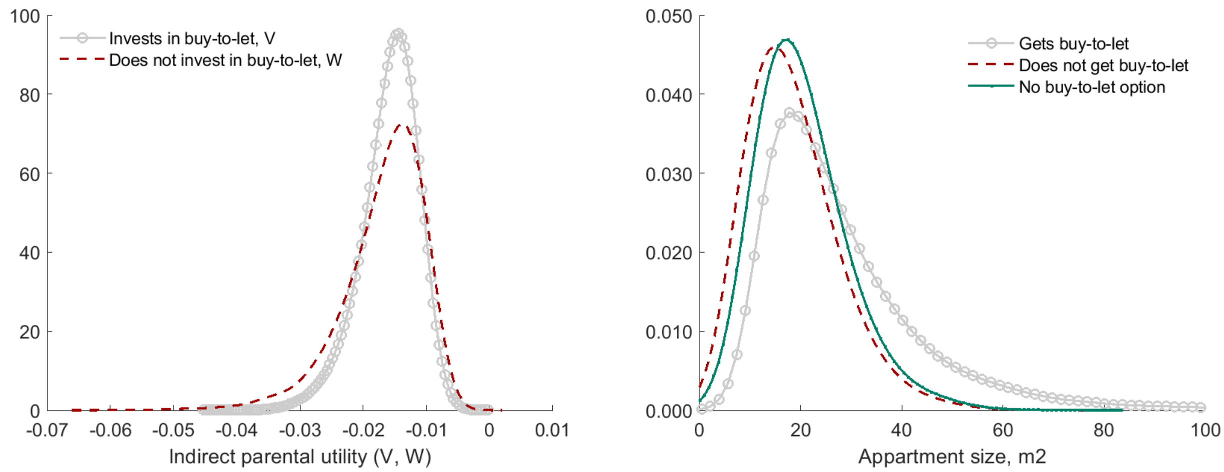

(a) Red dashed line: density of W for parents that do not invest in a buy-to-let. Grey circled line: density of V for parents who do invest in a buy-to-let.

(b) Green solid line: m^2 when all children must find rental on their own. Red dashed line: m^2 of children who do not get a buy-to-let. Grey circled line: m^2 of children in a buy-to-let.

Panel (b) of Figure 5 shows the housing outcomes for children whose parents invest in a buy-to-let (grey circled line), for those whose parents choose not to invest (red dashed) and for all children when the buy-to-let option is completely shut off (green). It appears that introducing the buy-to-let option allows children to live in larger apartments when their parents undertake the investment, as the right tail of this density is substantially thicker than the other two. To put in numbers, children in buy-to-let apartments live on average on 8 m^2 more than children in public rentals.

5.5 Counterfactual simulations

The real benefit of implementing a model for an economic decision process is that we can form a logically consistent assessment of the outcome if the setup changes. In this section, I consider how welfare and outcomes would change if the tax benefits for the buy-to-let arrangements were to be repealed. I also implement an equilibrium extension to the model in which the increased demand for housing through buy-to-let investments pushes up apartment prices. Below I first describe each scenario and then present the outcomes in conjunction for convenience.

Repealing the tax benefits

Repealing the tax benefits of the buy-to-let arrangement comes down to i) a removal of the option to retain rental income, ii) letting interest rate expenses from the buy-to-let figure directly into the budget constraint of the parent rather into the rental income equation and iii) requiring that the apartment is sold to the child at market value. In terms of the model setup, i) corresponds to fixing $w = 1$ and ii) means moving $\mathcal{R}(1-\gamma, P(h_k))$ from equation (18) into the LHS of first period parental budget constraint (19). For convenience, I obtain iii) by assuming that the down payment is disregarded, $\gamma = 0$, which then cancels the rebate offered to the child. The parent is therefore left with the choice variables $\{c_{p1}, c_{p2}, h_k\}$ in the optimization problem.

An equilibrium extension

As noted previously, the tax benefits of the buy-to-let arrangement work to subsidize such investments and they accordingly push up demand for housing. To illustrate the effect this has on housing prices, I evaluate the model subject to the condition that total housing demand must equal a fixed supply. The square meter price $\tilde{p}^m$ consistent with the minimization of excess demand is the equilibrium price. For this to work, I make two critical assumptions. The first is that the apartments inhabitated by children are situated in a market that is separate from the market on which parental housing is located. In other words, there is no feedback to the value of parental housing when the price of children housing increases. One may think of this as a geographical separation with children living close to study centers and parents elsewhere.

The second assumption is that the fixed supply of children housing equals the total

Table 8: Outcomes in counterfactual scenarios

	Baseline	**Repeal**	**Equilibrium**
Share investing in buy-to-let, %	67.09	59.21	70.60
Price pr m^2, DKK1000s	25.00	25.00	33.55
Average m^2, non buy-to-let tenants	20.31	20.26	14.98
Average m^2, buy-to-let tenants	28.30	26.36	20.77

Notes: Price is fixed during baseline evaluation and in repeal of buy-to-let tax benefits.

demand for housing by children in the case where no one is offered a buy-to-let. This amounts to assuming that the initial housing price p^m was clearing the children's housing market in absence of any buy-to-let investments. These assumptions make it a highly stylized example but note that I am mainly concerned with the qualitative effects.

Results

The effects of repealing the tax benefits and introducing equilibrium constraints on housing demand are are presented in Table 8 and Figure 6. Table 8 shows that repealing the tax benefits would decrease the share of parents investing in a buy-to-let by 8 percentage points. It also shows that the size of buy-to-let apartments fall by 2 m^2 on average when the benefits are withdrawn. These are not dramatic effects, but note that parents and children may still benefit from a buy-to-let investment even after tax benefits are repealed. The reason is that rental payments flow within the family rather than to a third party when investing, and because parents are still able to offer children a lower rent pr m^2 than obtains in the open market. Also, accounting for mortgage refinancing costs and stamp fees which lower the profitability of the arrangement would indeed make the repeal's effect more substantial.

An interesting effect of repealing the tax benefits is that it alters the distribution of apartment size for children in a buy-to-let. Comparing panel (d) of Figure 6 with (b) of Figure 5, we see that the distribution of buy-to-let sizes changes from having a comparatively thick right tail in the baseline to resembling that of the no buy-to-let scenario (green line).

In contrast to the repeal, imposing equilibrium on the market for children's housing

Figure 6: Distributions of welfare and apartment size in counterfactual scenarios

(a) **Equilibrium**: Density of welfare.

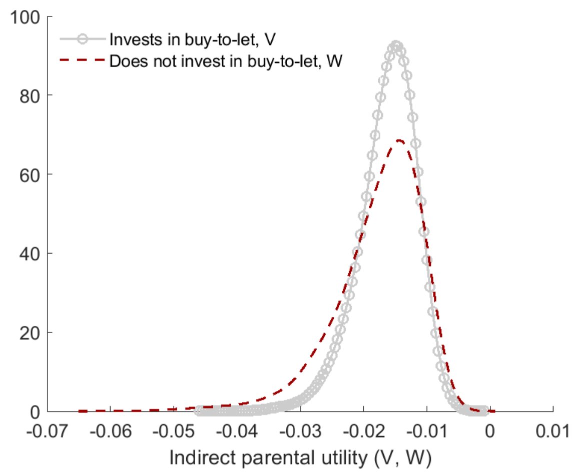

(b) **Repeal**: Density of welfare.

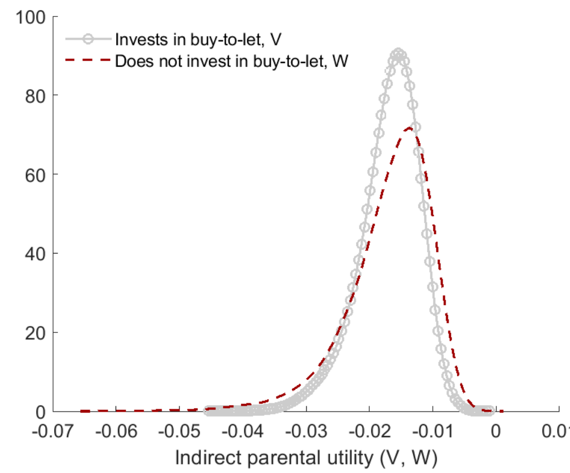

(c) **Equilibrium**: Density of apartment size, m^2.

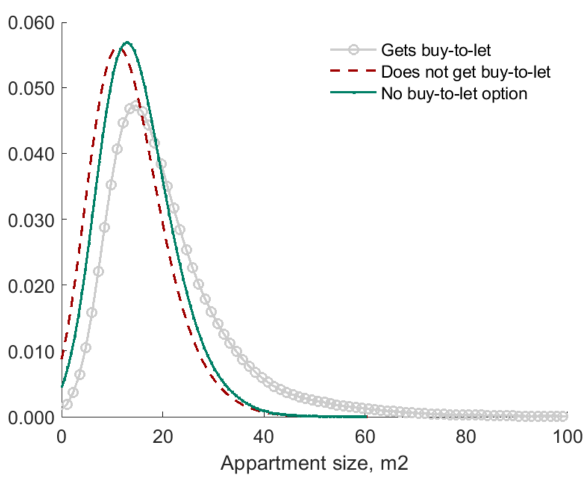

(d) **Repeal**: Density of apartment size, m^2.

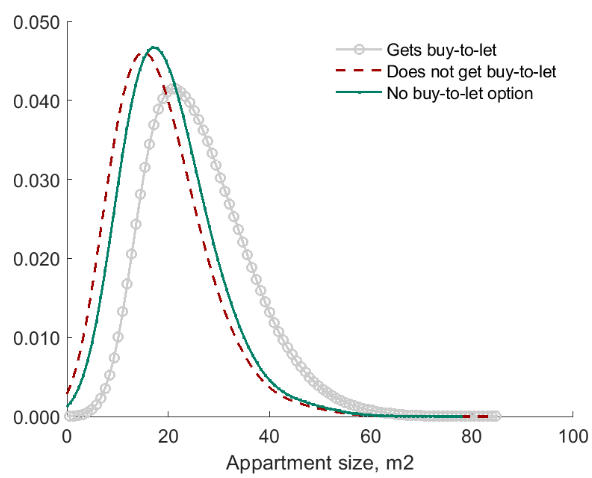

(a)-(b) Red dashed line: density of W for parents that do not invest in a buy-to-let. Grey circled line: density of V for parents who do invest in a buy-to-let.

(c)-(d) Green solid line: m^2 when all children must find rental on their own. Red dashed line: m^2 of children who do not get a buy-to-let. Grey circled line: m^2 of children in a buy-to-let.

does have rather dramatic effects. The price of square meters increases by 8,500 DKK and, interestingly, the share of parents investing in a buy-to-let also increases slightly. More parents invest in a buy-to-let as a reaction to the increased financial burden of servicing rental payments as it increases the utility of providing a discount on the market rent. In data, we do in fact observe that the number of buy-to-let investments closely track the development in housing prices. A more careful econometric analysis is needed to establish if this pattern is predominantly caused by a speculative motive of buy-to-let investors or because the option of providing affordable housing increases in value with housing prices. The simulation does however suggest an important effect stemming from the housing provision motive.

The fact that the equilibrium outcome implies increased prices hurts in particular children who are not provided a buy-to-let. Their average apartment size falls by 5.3 m^2, which keeps their monthly expenses roughly constant.

6 Conclusion

In this paper, I have considered the economic behavior around parental landlord arrangements in Denmark. The landlord arrangements are of special interest as they offer an opportunity to witness economic transfers among family members in full-population register data, which is otherwise difficult to come by. They also offer insights into the responsiveness to tax incentives since the arrangements are indirectly subsidized by a set of unintended tax benefits.

The empirical analysis shows that parents act in an altruistic and compensatory fashion when providing a rental for their children. In particular, the propensity to invest in a rental increases with parental income and decreases with contemporary child income. The same holds true for the price (hence quality) of the chosen rental. Interestingly, investment propensity and chosen apartment value increase with future child income, implying that parents help children smooth consumption over time through the rental investments. Furthermore, parents are observed to share capital gains on the rental investments with their children.

The theoretical framework for the parental landlord decision provided in the last part of the paper rationalizes the empirical findings. Future work will involve estimating this model directly on data. However, an externally calibrated version of the model indicates that repealing the tax benefits will discourage investments in parental rentals but not remove it all together.

References

Altonji, J. G., Hayashi, F., and Kotlikoff, L. J. (1997). Parental Altruism and Inter Vivos Transfers: Theory and Evidence. *Journal of political economy*, 105(6):1121–1166.

Andersen, A. L., Johannesen, N., and Sheridan, A. (2020). Bailing out the Kids: New Evidence on Informal Insurance from one Billion Bank Transfers. *SSRN Electronic Journal*.

Barczyk, D., Fahle, S., and Kredler, M. (2019). Save, Spend or Give? A Model of Housing, Family Insurance, and Savings in Old Age. *Mimeo*, page 94.

Barczyk, D. and Kredler, M. (2014). Altruistically motivated transfers under uncertainty: Altruistically motivated transfers under uncertainty. *Quantitative Economics*, 5(3):705–749.

Bernheim, B. D., Lemke, R. J., and Scholz, J. K. (2004). Do estate and gift taxes affect the timing of private transfers? *Journal of Public Economics*, 88(12):2617–2634.

Bernheim, B. D., Shleifer, A., and Summers, L. H. (1985). The strategic bequest motive. *Journal of Political Economy*, 93(5).

Boar, C. (2019). Dynastic Precautionary Savings. *Mimeo*, page 59.

Brown, M. (2006). Informal care and the division of end-of-life transfers. *Journal of Human Resources*, 41(1):191–219.

Browning, M., Chiappori, P.-A., and Weiss, Y. (2014). *Economics of the Family*. Cambridge Surveys of Economic Literature. Cambridge University Press.

Colliers (2019). Colliers Copenhagen Property Market Report. Technical report.

Davis, M. A. and Ortalo-Magné, F. (2011). Household expenditures, wages, rents. *Review of Economic Dynamics*, 14(2):248–261.

Escobar, S. and Ohlsson, H. (2019). Taxes, frictions and asset shifting - when Swedes disinherited themselves. *Mimeo*, page 61.

Hansen, J. Z. and Iversen, A. s. (2017). Ejerlejlighedslovens forbud mod omdannelse af udlejningslejligheder i småejendomme. Technical report, Danish Institute for Economic Modelling and Forecasting.

Hjelmar, U. (2009). Unge på boligmarkedet. Technical report, Anvendt KommunalForskning.

Hochguertel, S. and Ohlsson, H. (2009). Compensatory inter vivos gifts. *Journal of Applied Econometrics*, 24(6):993–1023.

Iman, R. L. and Conover, W.-J. (1982). A distribution-free approach to inducing rank correlation among input variables. *Communications in Statistics-Simulation and Computation*, 11(3):311–334.

Joulfaian, D. (2004). Gift taxes and lifetime transfers: time series evidence. *Journal of Public Economics*, 88(9):1917–1929.

Kaplan, G. (2012). Moving back home: Insurance against labor market risk. *Journal of Political Economy*, 120(3):446–512.

le Maire, D. and Schjerning, B. (2013). Tax bunching, income shifting and self-employment. *Journal of Public Economics*, 107:1–18.

Lockwood, L. M. (2018). Incidental Bequests and the Choice to Self-Insure Late-Life Risks. *American Economic Review*, 108(9):2513–2550.

McGarry, K. (2016). Dynamic aspects of family transfers. *Journal of Public Economics*, 137:1–13.

Municipality of Copenhagen (2019). Ungdomsboliganalyse til Kommuneplan. Technical report.

Rabe-Hesketh, S. and Skrondal, A. (2013). Avoiding biased versions of Wooldridge's simple solution to the initial conditions problem. *Economics Letters*, 120(2):346–349.

Sommer, E. (2017). Wealth Transfers and Tax Planning: Evidence for the German Bequest Tax. page 31.

Sommer, K. and Sullivan, P. (2018). Implications of US Tax Policy for House Prices, Rents, and Homeownership. *American Economic Review*, 108(2):241–274.

Train, K. E. (2009). *Discrete choice methods with simulation*. Cambridge university press.

Wooldridge, J. M. (2005). Simple solutions to the initial conditions problem in dynamic, nonlinear panel data models with unobserved heterogeneity. *Journal of Applied Econometrics*, 20(1):39–54. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jae.770.

Appendix A

Econometric model for discrete buy-to-let investment decision

As in Wooldridge (2005), $t = 0$ for the initial period of each child-parent(s) unit. btl_{it} denotes whether there was a buy-to-let investment within unit i in period t, $\mathbf{z}_{it}$ is the set of exogenous regressors and c_i is the family unit random effect. The general model is thus

$$btl_{it} = \mathbb{1}\{\mathbf{z}_{it}'\gamma + \rho btl_{it-1} + c_i + u_{it} > 0\},\ \ u_{it} \sim \mathcal{N}(0,1). \tag{31}$$

Following Rabe-Hesketh and Skrondal (2013), one may specify the distribution of c_i as a function of the average of time-varying regressors $\bar{\mathbf{z}}_i^+$ and their initial values $\mathbf{z}_{i0}$ together with the initial value of the outcome variable. Time-constant variables can only be used to model c_i and are therefore stored in $\bar{\mathbf{z}}_i^+$. Hence

$$c_i = \alpha_0 + \alpha_1 btl_{i0} + \bar{\mathbf{z}}_i^{+'}\boldsymbol{\alpha}_2 + \mathbf{z}_{i0}'\boldsymbol{\alpha}_3 + a_i, \tag{32}$$

where $a_i|(btl_{i0}, \bar{\mathbf{z}}_i^+, \mathbf{z}_{i0}) \sim \mathcal{N}(0, \sigma_a^2)$. One then substitutes (32) into (31) and uses that the density of btl_{it} given $(btl_{it-1}, \bar{\mathbf{z}}_i^+, \mathbf{z}_{i0}, a_i)$ follows a probit. Integrating out the random effect and letting $(btl_{i0} = btl_0, \bar{\mathbf{z}}_i^+ = \bar{\mathbf{z}}^+, \mathbf{z}_{i0} = \mathbf{z}_0)$ yields the density for a cross-sectional unit outcome $(btl_1, ..., btl_T)$ as

$$\int_{\mathbb{R}} \prod_{t=1}^{T} \{\Phi(\mathbf{z}_t'\gamma + \rho btl_{t-1} + \alpha_0 + \alpha_1 btl_{i0} + \bar{\mathbf{z}}^{+'}\boldsymbol{\alpha}_2 + \mathbf{z}_0'\boldsymbol{\alpha}_3)^{btl_t} \times \\ [1 - \Phi(\mathbf{z}_t'\gamma + \rho btl_{t-1} + \alpha_0 + \alpha_1 btl_0 + \bar{\mathbf{z}}^{+'}\boldsymbol{\alpha}_2 + \mathbf{z}_0'\boldsymbol{\alpha}_3)]^{(1-btl_t)}\}(1/\sigma_a)\phi(a/\sigma_a)da \tag{33}$$

Additional regressions and tables for Section 4

Table 9: Dynamic Random Effects probit of buy-to-let decision. Year dummies and intial conditions for regression in Table 4.

	Coefficients	S.E.
Buy-to-let$_{t0}$	0.8169***	0.0247
Parental income$_{t_0}$	0.0088***	0.0007
Child income$_{t_0}$	-0.0467***	0.0050
$\mathbb{1}\{$Child in high educ$_{t_0}\}$	-0.1541***	0.0134
$\mathbb{1}\{$Child age $\geq 24_{t_0}\}$	-0.0251	0.0190
$\mathbb{1}\{$Self-employed parent$_{t_0}\}$	0.0002	0.0126
$\mathbb{1}\{z^* <$ Parental income$_{t_0}\}$	0.0796***	0.0079
$\mathbb{1}\{$Child has graduated$_{t_0}\}$	0.1896***	0.0389
$\mathbb{1}\{t = 1996\}$	0.0410**	0.0176
$\mathbb{1}\{t = 1997\}$	0.1022***	0.0169
$\mathbb{1}\{t = 1998\}$	0.2029***	0.0162
$\mathbb{1}\{t = 1999\}$	0.2427***	0.0160
$\mathbb{1}\{t = 2000\}$	0.2150***	0.0161
$\mathbb{1}\{t = 2001\}$	0.1896***	0.0162
$\mathbb{1}\{t = 2002\}$	0.1971***	0.0163
$\mathbb{1}\{t = 2003\}$	0.2336***	0.0163
$\mathbb{1}\{t = 2004\}$	0.2072***	0.0165
$\mathbb{1}\{t = 2005\}$	0.2952***	0.0164
$\mathbb{1}\{t = 2006\}$	0.2224***	0.0169
$\mathbb{1}\{t = 2007\}$	0.1885***	0.0172
$\mathbb{1}\{t = 2008\}$	0.1128***	0.0175
$\mathbb{1}\{t = 2009\}$	0.2788***	0.0177
$\mathbb{1}\{t = 2010\}$	0.1022***	0.0181
$\mathbb{1}\{t = 2011\}$	0.1359***	0.0184
$\mathbb{1}\{t = 2012\}$	0.0917***	0.0192
$\mathbb{1}\{t = 2013\}$	0.3438***	0.0196
$\mathbb{1}\{t = 2014\}$	0.3198***	0.0209
$\mathbb{1}\{t = 2015\}$	0.33***	0.02

Table 10: Dynamic random effects probit of buy-to-let investment decision, only home-owning parents.

	Coefficients	S.E.
Constant	-3.2774***	0.0263
Buy-to-let$_{t-1}$	2.8615***	0.0098
Parental income$_t$	0.0362***	0.0021
Child income$_t$	-0.0544***	0.0029
$\mathbb{1}\{$Child in high educ$_t\}$	0.3871***	0.0067
$\mathbb{1}\{$Child age $\geq 24_t\}$	-0.2057***	0.0062
$\mathbb{1}\{$Self-employed parent$_t\}$	0.0145	0.0143
$\mathbb{1}\{z^* <$ Parental income$_t\}$	-0.0374***	0.0119
Parental income$_t \times \mathbb{1}\{z^* <$ Parental income$_t\}$	-0.0044**	0.0022
$\mathbb{1}\{$Child has graduated$_t\}$	-0.1289***	0.0095
$\mathbb{1}\{$Parent is divorced$\}$	-0.2993***	0.0062
N properties owned	0.1357***	0.0024
N siblings	-0.0914***	0.0031
$\mathbb{1}\{$Parent w. high education$\}$	0.1085***	0.0057
$\mathbb{1}\{$Child does graduate$\}$	0.0085	0.0077
E$_i$[Parental income]	0.0008	0.0015
E$_i$[Child income]	-0.0180***	0.0044
E$_i[\mathbb{1}\{$Child in high educ$\}]$	-0.1295***	0.0159
E$_i[\mathbb{1}\{$Child age $\geq 24\}]$	0.1494***	0.0237
E$_i[\mathbb{1}\{$Self-employed parent$\}]$	0.1455***	0.0204
E$_i[\mathbb{1}\{z^* <$ Parental income$\}]$	0.2366***	0.0170
E$_i[$Parental income$\times\mathbb{1}\{z^* <$ Parental income$\}]$	0.0093***	0.0023
E$_i[\mathbb{1}\{$Child has graduated$\}]$	-0.2004***	0.0220
σ_a	0.5357	-
Log-likelihood	-216953.36	-
N	5187688	-
N_i	587141	-

*Notes: *** significant at 1% level, ** at 5%, * at 10%. Prices and incomes are in DKK 2015 and scaled by 100,000. Z-scores in parentheses. z^* denotes the cut-off between mid or bottom taxation (depending on the period). $E_i[v]$ denotes the average of variable v for cross-sectional unit i. Initial conditions and year dummies are omitted from the table.*

Table 11: Dynamic random effects probit of buy-to-let investment decision, only home-owning parents with child living in Copenhagen for at least one year between ages 18-30.

	Coefficients	S.E.
Constant	-3.1479***	0.0321
Buy-to-let$_{t-1}$	2.8321***	0.0115
Child income$_t$	-0.0495***	0.0034
Parental income$_t$	0.0337***	0.0024
$\mathbb{1}\{z^* <$ Parental income$_t\}$	-0.0476***	0.0143
Parental income$_t \times \mathbb{1}\{z^* <$ Parental income$_t\}$	-0.0024	0.0026
$\mathbb{1}\{$Child in high educ$_t\}$	0.3546***	0.0081
$\mathbb{1}\{$Child age $\geq 24_t\}$	-0.1888***	0.0074
$\mathbb{1}\{$Self-employed parent$_t\}$	0.0264	0.0173
$\mathbb{1}\{$Child has graduated$_t\}$	-0.1378***	0.0115
$\mathbb{1}\{$Parent is divorced$\}$	-0.3024***	0.0073
$\mathbb{1}\{$Parent w. high education$\}$	0.0950***	0.0070
$\mathbb{1}\{$Child does graduate$\}$	0.0103	0.0093
N properties owned	0.1355***	0.0030
N siblings	-0.0864***	0.0037
E$_i$[Child income]	-0.0212***	0.0050
E$_i$[Parental income]	-0.0002	0.0017
E$_i[\mathbb{1}\{z^* <$ Parental income$\}]$	0.2080***	0.0207
E$_i$[Parental income$\times \mathbb{1}\{z^* <$ Parental income$\}]$	0.0092***	0.0027
E$_i[\mathbb{1}\{$Child in high educ$\}]$	-0.1676***	0.0190
E$_i[\mathbb{1}\{$Child age $\geq 24\}]$	0.0154	0.0293
E$_i[\mathbb{1}\{$Self-employed parent$\}]$	0.1517***	0.0248
E$_i[\mathbb{1}\{$Child has graduated$\}]$	-0.1721***	0.0272
σ_a	0.5192	-
Log-likelihood	-148060.74	-
N	3027525	-
N_i	337477	-

*Notes: *** significant at 1% level, ** at 5%, * at 10%. Prices and incomes are in DKK 2015 and scaled by 100,000. Z-scores in parentheses. z^* denotes the cut-off between mid or bottom taxation (depending on the period). $E_i[v]$ denotes the average of variable v for cross-sectional unit i. Initial conditions and year dummies are omitted from the table.*

Table 12: Year dummies for linear regression of buy-to-let purchase price, see Tabel 5 for remaining estimates.

	ln(Price)		
	Copenhagen	Aarhus	Elsewhere
$\mathbb{1}\{t = 1995\}$	-0.0049	0.0398	-0.0105
$\mathbb{1}\{t = 1996\}$	0.0771***	0.1409***	0.0394
$\mathbb{1}\{t = 1997\}$	0.1668***	0.2325***	0.1020***
$\mathbb{1}\{t = 1998\}$	0.2922***	0.3447***	0.2332***
$\mathbb{1}\{t = 1999\}$	0.4577***	0.3802***	0.2624***
$\mathbb{1}\{t = 2000\}$	0.5849***	0.4079***	0.3452***
$\mathbb{1}\{t = 2001\}$	0.7766***	0.4305***	0.3771***
$\mathbb{1}\{t = 2002\}$	0.8379***	0.4821***	0.4244***
$\mathbb{1}\{t = 2003\}$	0.9155***	0.5427***	0.4659***
$\mathbb{1}\{t = 2004\}$	1.0129***	0.7151***	0.5626***
$\mathbb{1}\{t = 2005\}$	1.1977***	0.9059***	0.7298***
$\mathbb{1}\{t = 2006\}$	1.3915***	0.9721***	0.7973***
$\mathbb{1}\{t = 2007\}$	1.2512***	0.9304***	0.8312***
$\mathbb{1}\{t = 2008\}$	1.1249***	0.8346***	0.7518***
$\mathbb{1}\{t = 2009\}$	0.9694***	0.7788***	0.6396***
$\mathbb{1}\{t = 2010\}$	1.0079***	0.7975***	0.6606***
$\mathbb{1}\{t = 2011\}$	0.9604***	0.8034***	0.5557***
$\mathbb{1}\{t = 2012\}$	0.8675***	0.7392***	0.5701***
$\mathbb{1}\{t = 2013\}$	1.0025***	0.7513***	0.5079***
$\mathbb{1}\{t = 2014\}$	1.0543***	0.9006***	0.5868***
$\mathbb{1}\{t = 2015\}$	1.1637***	0.7073***	0.9098***

*Notes: *** significant at 1% level, ** at 5%, * at 10%. Based on robust Z-scores.*

Table 13: Buy-to-let apartments sold by parents

Year	Sells to third party	Sells to child
1999	434	136
2000	504	168
2001	525	164
2002	593	180
2003	650	244
2004	765	335
2005	1013	720
2006	678	464
2007	484	325
2008	343	189
2009	246	45
2010	374	144
2011	360	107
2012	332	79
2013	421	123
2014	588	143

Appendix B Distributions of data for numerical model

Approximated joint distribution of state variables

I use Latin hypercube sampling (LHS) in order to construct a joint distribution of all state variables that enter into the collective model. This requires a target correlation matrix of the state variables, provided in Table 15, and their independent distributions, see Table 14. The LHS technique is an efficient way of drawing from a distribution as it uses full coverage of the distribution's support. Using LHS to create a multivariate sample of N observations from K correlated variables proceeds as follows. The cumulative distribution of each variable k is stratified into N equiprobable intervals. One random value is drawn from each interval of every distribution, hence NK values in total, and then transformed using the inverse CDF. The sampled values are now reordered to have the same rank order as prescribed by the target (linear) correlation structure. This step is solved for using Choleski decomposition of the target correlation matrix and the sample matrix, see Iman and Conover (1982) for details.

Table 14: Fitted of distributions of state variables

Description	**Parameter**		**Distribution**
Parent income period 1	z_{p1}	$\sim$	GEV(311.5, 162.3, 0.012)
Parent income period 2	z_{p2}	$\sim$	GEV(314.9, 161.2, 0.025)
Child income period 1	z_{k1}	$\sim$	GEV(99.8, 55.6, -0.036)
Child income period 2	z_{k2}	$\sim$	GEV(112.2, 56.9, 0.002)
Size of parent's home in m^2	h_p	$\sim$	GEV(126.2, 38.3, -0.116)
Price pr m^2 for parents	p_p^m	$\sim$	Piecewise linear approximation.*
Initial LTV of parents	λ_{p0}	$\sim$	Piecewise linear approximation.*

**See Appendix Figure 7 for visualization of distributions.*

Table 15: Cross correlations between state variables

	Parental income$_t$	Parental income$_{t-1}$	Kid income$_t$	Kid income$_{t-1}$	Parental LTV$_t$	Home m2, parents$_t$	Price pr m2, parents$_t$
Parental income$_t$	1.000	-	-	-	-	-	-
Parental income$_{t-1}$	0.884	1.000	-	-	-	-	-
Kid income$_t$	-0.057	-0.061	1.000	-	-	-	-
Kid income$_{t-1}$	-0.061	-0.064	0.701	1.000	-	-	-
Parental LTV$_t$	-0.102	-0.106	0.054	0.058	1.000	-	-
Home m2, parents$_t$	0.206	0.203	-0.053	-0.048	-0.069	1.000	-
Price pr m2, parents$_t$	0.280	0.288	-0.024	-0.032	-0.234	-0.139	1.000

Figure 7: Distributions of model state variables, kids at age 22 in 2012

(a) Density of parents' home size

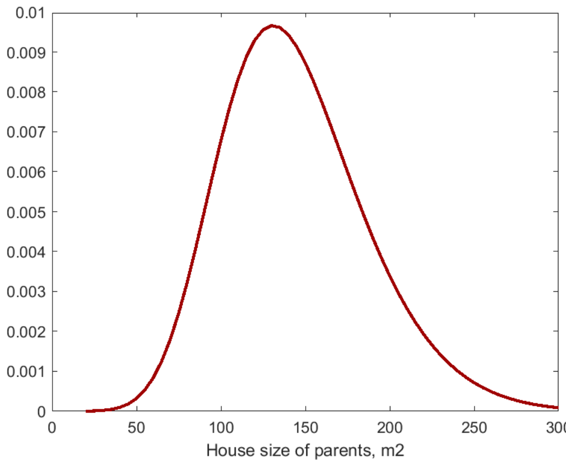

(b) Density price pr m2 for parents

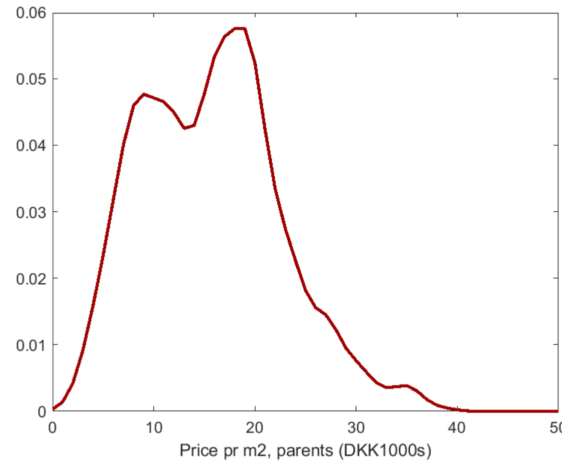

(c) Kernel density of parents' LTV

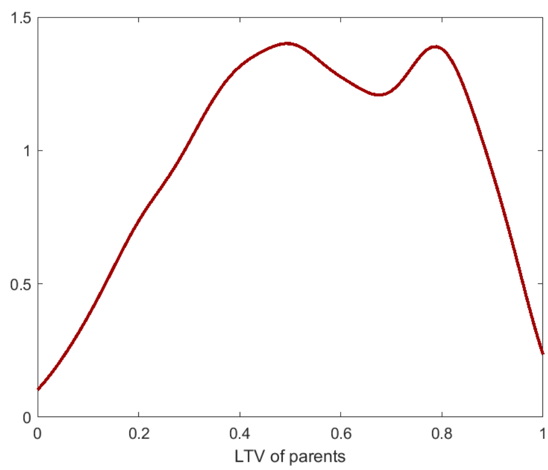

(d) Income distributions of parents and children.

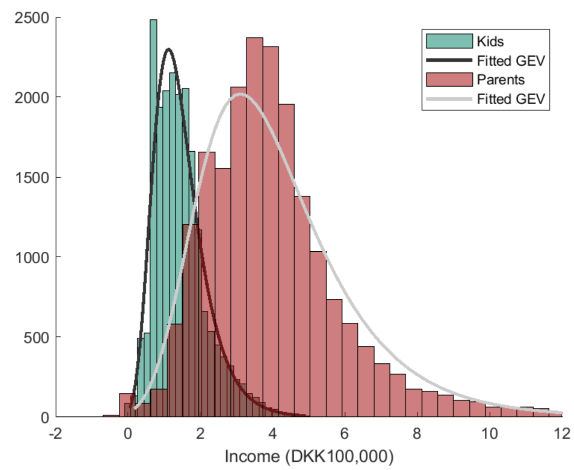

Notes: Only parent-children pairs where i) the child participates in further education between age 18-27 and ii) the child lives in the Copenhagen Metro Area for at least 1 year between age 18-27. Incomes deflated to 2015 prices.

Figure 8: Annual rent for buy-to-lets and rentals on the open market.

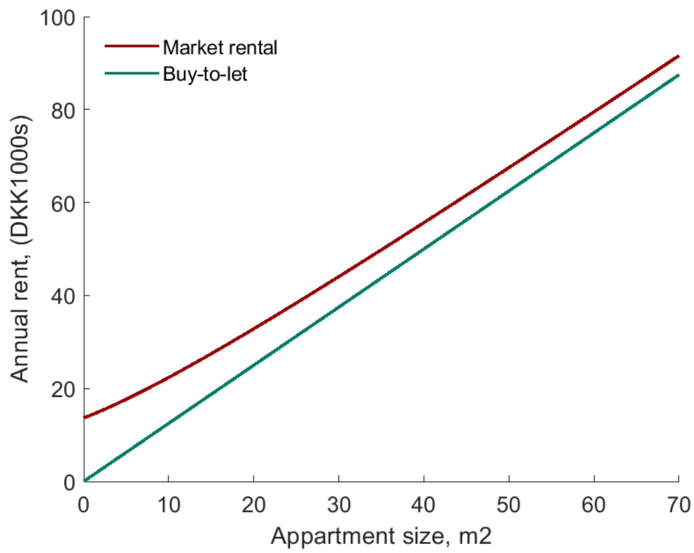

Notes: Based on parameters in Table 7.