Build, Baby, Build: How Housing Shapes Fertility

Benjamin K. Couillard*

Economics Department, University of Toronto

November 2, 2025

Click Here For Latest Version

Abstract

Many developed countries face low and falling birthrates, potentially affected by rising costs of housing. Existing evidence on the fertility-housing cost relationship typically uses geographic variation (raising selection issues), neglects unit size, and says little about policy. To progress on these fronts, I first specify a dynamic model of the joint housing-fertility choice allowing choices over location and house size, estimated using US Census Bureau data. I extend 'micro-moment' techniques (Petrin, 2002; Berry *et al.*, 2004a) both to circumvent data constraints and to incorporate heterogeneous residuals, which can prevent misspecification. Housing choice estimates confirm a Becker quantity-quality model's predictions: large families are more cost-sensitive, and so rising housing costs disincentivize fertility. To study the causal effect of rising housing costs on fertility, I vary them directly within the model, finding that rising costs since 1990 are responsible for 11% fewer children, 51% of the total fertility rate decline between the 2000s and 2010s, and 7 percentage points fewer young families in the 2010s. Policy counterfactuals indicate that a supply shift for large units generates 2.3 times more births than an equal-cost shift for small units: family-friendly housing is the more important policy lever.

JEL: J13, R21, R52, J12, R312, C81

Keywords: Fertility decline, housing affordability, household formation, family-friendly housing, neighborhood sorting, housing supply, zoning/regulatory tax, dynamic discrete choice, micro-moments, iterative proportional fitting

*University of Toronto Economics Department. Email: ben.couillard@mail.utoronto.ca. I thank my supervisory committee, Nate Baum-Snow, Stephan Heblich, Rob McMillan, & Victor Aguirregabiria, for their steadfast guidance and support. I also thank U of T's urban and IO groups, and participants of the 2025 Canadian Summer Conference in Real Estate and Urban Economics. This research was supported in part by funding from the Social Sciences and Humanities Research Council of Canada, and an Ontario Graduate Scholarship. All errors and omissions are my own.

1 Introduction

Persistently low fertility has a number of adverse economic consequences, including higher pension and healthcare costs, less innovation, slower growth, and population decline (Bloom *et al.*, 2024). Immigration is not a durable solution since two thirds of the global population live in a country with below replacement-rate fertility and the global fertility rate will continue declining through 2100 (UN, 2022), and so the onus is on governments to raise birth rates. A promising but under-studied avenue is housing: larger families demand more housing and so a higher marginal cost may deter a marginal birth, just as in a quantity-quality model (Becker, 1960) a high cost of child quality (*e.g.*, education) may deter a marginal birth. Empirical support for this channel exists, but is often afflicted by a geographic selection problem as those with preferences for many children sort into lower-cost locations, is unable to accommodate nonlinearities in the cost of housing over unit size, and does not consider the impacts of policies that shift housing supply.

I build on this reduced-form literature by specifying a structural model of dynamic joint housing-fertility and household formation choice, which allows me to obtain the causal impact of housing costs on fertility by varying them directly, and which permits general equilibrium counterfactuals targeting particular unit sizes. The model nests a static residential choice model in a broader dynamic living arrangement model, which includes fertility choice. The residential choice model specifies the choice variable as combinations of neighborhood and number of bedrooms, and allows preferences to be heterogeneous over age, tenure, family/roommate households, and household size. Agents make living arrangement and fertility decisions knowing how they will affect their housing demand, and so changes in housing characteristics (including costs) affect the values of living arrangement choices. The living arrangement decision includes living with parents or other family members, and so the long-run value of having children is linked to the endogenous probability that adult children or other family members will unilaterally choose to live in the household. This mechanism is important when units are heterogeneous in size, since affordable small units raise the short-run value of being single relative to starting a family, but also reduces the long-run cost of fertility since adult children are more likely to move out.

I estimate the housing demand model by applying and extending IO "micro-moment" techniques to surmount the problem of not observing the joint distribution of household counts over census tracts and the many dimensions of demographic heterogeneity. Since Petrin (2002) and Berry *et al.* (2004a), marginal distributions of the choice variable with individual dimensions of heterogeneity have been used to estimate heterogeneous parameters in discrete choice models under an assumption of additive separability. However, these techniques restrict the utility residual to be homogeneous over demographics and can thus lead to model misspecification in some cases, depending on which marginal distributions are used, the number of choice options, and the number of parameters in the

model. I show that in the case where the utility residual is unrestricted and there are no random coefficients, the observed marginal distributions are sufficient statistics for the two-step estimator of Bayer *et al.* (2007).

Housing demand results are consistent with a Beckerian (1960) quantity-quality model extended to include housing: larger families are more sensitive to housing costs and have stronger tastes for larger units. Thus, rising housing costs – in only the largest units, or on average – do contribute to declining fertility. I quantify this effect in a decomposition, using fertility utility parameters to hold constant all non-housing determinants of fertility: if rents had not risen since 1990, 13 million (11%) more children would have been born between 1990 and 2020, and the decline in fertility between the 2000s and 2010s would have been smaller by 51%. In further decompositions, I separately impose 1990 relative rents (letting the average rents rise) and average rents (letting the relative rents vary), finding that average rents have the larger effect on fertility.

I then compare two potential housing supply policy counterfactuals, separately shifting the supply of large 3+ bedroom and small 1 bedroom units. I interpret these shifts as reductions in distortionary local regulatory taxes, however to put the two policies on equal footing I define the shifts as if subsidies that generate equal aggregate funding had been provided (5% of annual aggregate rental expenditures, assuming all agents are renters and based on baseline rents and quantities), so that the small-unit shift generates more total housing units. The small unit policy is *de facto* YIMBY policy, which focuses on small and low-cost units, while the large unit policy is a direct attempt at family-friendly housing policy. In principle, the YIMBY policy could cause a larger fertiliy increase with both generate a larger effect on aggregate rents and reduce the long-run housing cost of fertility by prompting young adults to move out of their parents' homes, even as it would steepen the size-cost gradient. On the other hand, the large unit policy would generate a smaller decrease in average rents, but flatten this gradient. I find that the large unit policy causes 4.7 million (4%) more children to be born over three decades. Interpreting these shifts as a subsidy, the return is one birth per $160,000 (2022 USD), which is on the high end of estimates; interpreting these shifts as reductions in "regulatory taxes" on development suggests that birth rates can be increased at a low cost while increasing economic efficiency. The YIMBY policy achieves only 43% this number of births, because its largest effect is making living alone more attractive.

I conclude that rising housing costs are a major cause of fertility decline, even after correcting for endogenous sorting. This relationship mostly results from increasing housing costs across all unit sizes, rather than raising the relative rents of large units. However, a large unit policy that combines a smaller effect on average rents with a flattening size-cost gradient outperforms a small unit policy that combines a larger effect on average rents with a steepening size-cost gradient. If housing is to be a lever in family policy, the focus must be on producing the housing that families want.

1.1 Related Literature and Contribution

My first contribution is to provide estimates of the impact of housing costs on fertility that are free of selection bias. Although the negative relationship between income and fertility is "pervasive," (Doepke *et al.*, 2023), it is generated by other factors like a quantity-quality tradeoff (Becker, 1960) and opportunity cost of women's time. Indeed, a host of empirical papers show that children are a normal good (Fleisher and Rhodes, 1979; Borg, 1989; Heckman and Walker, 1990; Black *et al.*, 2013; Kearney and Wilson, 2018; Autor *et al.*, 2019; Cesarini *et al.*, 2023; Tsai *et al.*, 2022). When housing costs are high, there is less income available for children, and when the relative price of family-friendly housing is high then families will substitute towards fewer children, although incumbent homeowners may experience a countervailing positive wealth shock even as mortgage payments increase.

The empirical literature on housing costs and fertility is consistent with these mechanisms, finding that high housing costs reduce renter fertility and have varying effects on owner fertility (Dettling and Kearney, 2014; Clark and Ferrer, 2019; Lovenheim and Mumford, 2013; Pan and Yang, 2022; Daysal *et al.*, 2021; Atalay *et al.*, 2021; Liu and Zhang, 2024; Clark, 2012; Simon and Tamura, 2009; Dettling and Kearney, 2025; Cumming and Dettling, 2024; Fazio *et al.*, 2025). However, all of the aforementioned papers save Cumming and Dettling (2024) and Fazio *et al.* (2025) (which focus on housing finance) use geographic variation in housing costs, and provided that households simultaneously choose both fertility and housing reduced-form coefficients will include a sorting effect as high-fertility-propensity households sort into lower-cost locations. I explicitly account for this phenomenon by embedding a sorting model into a fertility choice model, and provide estimates of the impact of rising housing costs since 1990 in the first decomposition, generating smaller aggregate results than an analogous reduced-form exercise.

My second contribution is to analyze the distinct roles that small and large units play in fertility and household formation by unifying them into a single framework. Young adults rarely start families while living with their parents. An empirical literature finds that high housing costs reduce household formation (Börsch-Supan, 1986; Haurin *et al.*, 1993; Ermisch and Salvo, 1997; Ermisch, 1999; Paciorek, 2013; Cooper and Luengo-Prado, 2018; Wrenn *et al.*, 2019; Lafortune and Low, 2023) but both this literature and the housing-fertility literature (with the exception of Clark, 2012) do not consider housing unit size. We would expect the affordability of large units to be more relevant for fertility and the affordability of small units to be more relevant for household formation. Affordable small units have an ambiguous effect on fertility: they may lower the long-run cost of fertility by inducing adult children to move out, or may make remaining childless more attractive relative to starting a family[1]. I accommodate these mechanisms by including the number of bedrooms

[1]They may also provide young adults an intermediate step to starting a family by fostering independence and attractiveness on the marriage market, but this mechanism is not included in the model.

as a dimension of housing choice, and including the presence of adult children and other family members in the dynamic state. These mechanisms are explored in the latter decompositions where I alternately hold constant and let vary average and relative housing prices, finding that average prices matter more, as well as the housing policy counterfactuals.

My third contribution is to analyze fertility under housing policy counterfactuals. The large number of reduced-form papers on the topics of housing costs, fertility, and household formation underlines the importance of these topics, but housing costs are endogenous and cannot be directly controlled by policy, calling for counterfactual analysis that shifts exogenous variables. There is increasing interest in using structural models to jointly analyze urban economics and demography, but these papers focus more on the impact of demographics on urban outcomes rather than the reverse (Coeurdacier *et al.*, 2022; Ahlfeldt *et al.*, 2025; Moreno-Maldonado and Santamaría, 2024; Albouy and Faberman, 2025). I therefore consider housing supply subsidies (that can also be understood as a reduction in regulatory tax) in conjunction with estimated local supply elasticities (Baum-Snow and Han, 2024), solve for the new sorting-fertility equilibrium, and calculate fertility statistics for comparison, finding that subsidizing larger units is far more effective in raising fertility.

My fourth contribution is to include rich heterogeneity *and* microgeography in a neighborhood sorting model *without* confidential microdata by applying IO micro-moment techniques (Petrin, 2002; Berry *et al.*, 2004a), *and* deriving a sufficient statistics result to extend them to to accommodate heterogeneous residuals, which can prevent misspecification in some cases. In comparison, logit-based sorting models have been estimated with rich heterogeneity and microgeography using confidential microdata (*e.g.* Bayer *et al.*, 2007); with rich heterogeneity and larger geographies observable in public-use microdata (*e.g.* Diamond, 2016); and with limited heterogeneity and microgeography using public-use tabulations (*e.g.* Couture and Handbury, 2020). Instead, I employ a simple manipulation of the Bayer *et al.* (2007) log-likelihood under additive separability of preferences to show that public-use tabulations and higher-geography public-use microdata are sufficient. I pursue this approach because standard IO tools that use the same data (like PyBLP, Conlon and Gortmaker, 2020; Conlon and Gortmaker, 2023) impose demographic homogeneity on the systematic unobserved utility ξ. As Athey and Imbens (2007) first observed, a model with a single unobserved characteristic may not be able to rationalize choice data that is disaggregated by demographics. I show that when the number of micro-moments (marginal distributions) grows with the number of options J, there is a J^* such that a model with a homogeneous residual is unable to rationalize the data and is thus misspecified, and standard identification and consistency results (Berry and Haile, 2014; Berry *et al.*, 2004b) may not hold. However, identification does hold when the utility residual varies over demographics (Berry and Haile, 2010) and consistency very likely extends to this case as well.

1.2 Outline

In Section 2, I introduce the data and provide descriptive evidence. In Section 3, I introduce identifying variation, replicate reduced-form findings, and discuss the impact of sorting. In Section 4, motivated by descriptive evidence, limitations of the reduced-form, and a quantity-quality model (Becker and Lewis, 1973) extended to include housing, I introduce the structural model. In Section 5, I discuss estimation of the structural model. In Section 6, I define the partial and general equilibrium concepts used to solve decompositions and counterfactual scenarios respectively. In Section 7, I present the results of estimation, decompositions, and counterfactuals. Section 8 concludes.

2 Data and Descriptive Statistics

2.1 Data

All data except for that used to construct one of the instruments comes from the 1990 or 2000 US Decennial Census or 2008-12 or 2018-22 5-year American Community Survey, by way of IPUMS USA (Ruggles *et al.*, 2025) or NHGIS (Schroeder *et al.*, 2025).

Microdata. The fertility part of the model uses microdata and ignores its limited geographic information. Because fertility is part of a broader living arrangement choice, the most important information is related to household structure. Although the Census and ACS are cross-sectional and not longitudinal, I use information on the ages of children in the household, the duration of time spent in their current residence, and the presence of a mortgage to infer the household structure and tenure state a decade prior. Because the US Census, like censuses across the world, under-counts young children (U.S. Census Bureau, 2016), I re-weight the estimation sample so that it matches age-specific decadal fertility rates based on official estimates from the US CDC (Martin *et al.*, 2013; Driscoll and Hamilton, 2025), and thus the official US total fertility rate, using the iterative proportional fitting (IPF) algorithm. I also use microdata to construct the supplementary city-level geographic tabulation referred to immediately below.

Geographic Tabulations. Housing demand estimation would be straightforward if the joint distribution of households over census tract, number of bedrooms in the housing unit, and demographics was observed. Unfortunately, it is not observable outside of US Census Bureau Research Data Centers. Past papers in the literature have used RDC data, compromised on geography, or compromised on demographic heterogeneity. Instead, I use a set of marginal distributions of household counts over combinations of census tract (or the city it belongs to), number of bedrooms in the house, and dimensions of demographic heterogeneity, which are *sufficient statistics* for the parameters of the model. One of these marginal distributions is over *all* dimensions of demographic heterogeneity considered, and the *city* that a census tract belongs to. I combine these marginal

distributions into a single rectangular dataset using iterative proportional fitting (IPF). This estimate of the joint distribution exactly matches the input marginals, and can thus be conveniently used in a computational routine requiring a rectangular dataset for an estimator for which the marginal distributions are sufficient statistics.

To better illustrate the data used, consider the following simple example (tables shown in Section A1. There are four census tracts – each of which belongs to one of two cities – and two age groups and two sexes. This is the unobserved joint distribution, depicted in Table 5's "Population" column. Because this joint distribution is unobserved, we would instead use three marginal distributions in Table 6: tract by age, tract by sex, and age by sex by city. These marginal distributions could be used directly in a bespoke maximum likelihood estimation routine, or they could first be combined into a rectangular estimate of the joint distribution using IPF (Table 5's right-most column) to be passed into a more standard maximum likelihood estimation routine.

Explanatory variables used in housing demand estimation and data for reduced-form analysis – rents, share single-family homes, fertility rates (Section 3.1) – are also constructed using census tract tabulations.

Rents. I use rents to measure housing costs, which as a flow cost are more conceptually appealing than house prices while being highly correlated with house prices and mortgage payments. The Census Bureau produces tabulations of occupied housing units by census tract and nominal rent bins; from these I estimate the median nominal rent using a linear cumulative density function and convert to real terms, and I scale as necessary to match year-bedroom averages in the microdata.

Railway Noise. Noise pollution from railways is used to construct one of two instruments for rents used in this paper; a discussion of identifying assumptions is deferred. The US Department of Transportation published the average daily noise in decibels from trains in 2018 over a fine grid. This measure is calculated based on a noise model that takes train microdata as an input. Data is coded as missing if the noise is less than 45 decibels (roughly ambient background noise), which I replace with 40 decibels.

Consolidated Census Tracts. Because census tracts are redrawn for each census, I construct a balanced panel of about 48,000 neighborhoods by minimally aggregating census tracts to create consistent units[2]. That is, if a 1990 tract is split into two for the 2000 census, I take the 1990 tract as the unit and aggregate as necessary in later years to generate counts for that unit.

[2]Due to small imperfections in the geospatial data representing census tracts, there are a number of tiny splits when comparing tract vintages that are in fact meaningless. I ignore splits that imply a tract is crossing a county border (which is impossible), is home to zero people (based on a spatial merge of census *block* population data as points), comprise less than 5% of the area of any of the units involved, or which cross Public Use Microdata Area (PUMA) boundaries (which is important for a technical estimation reason discussed in Section A9.2).

2.2 Descriptive Statistics

Housing and Fertility are Jointly Chosen. Anecdotal evidence abounds for the idea that housing and fertility are jointly chosen: young couples often time home purchases and moves with starting a family. In Figure 1, I show the results of a regression of a birth indicator on housing characteristics that captures the strength of this correlation. When all housing characteristics are moved from the non-family-friendly value to the family-friendly value, the conditional probability of giving birth increases from less than 1% to nearly 20%, compared with a mean of 11%. This non-causal relationship is economically significant, and so we could learn a lot from a framework that respects this joint determination.

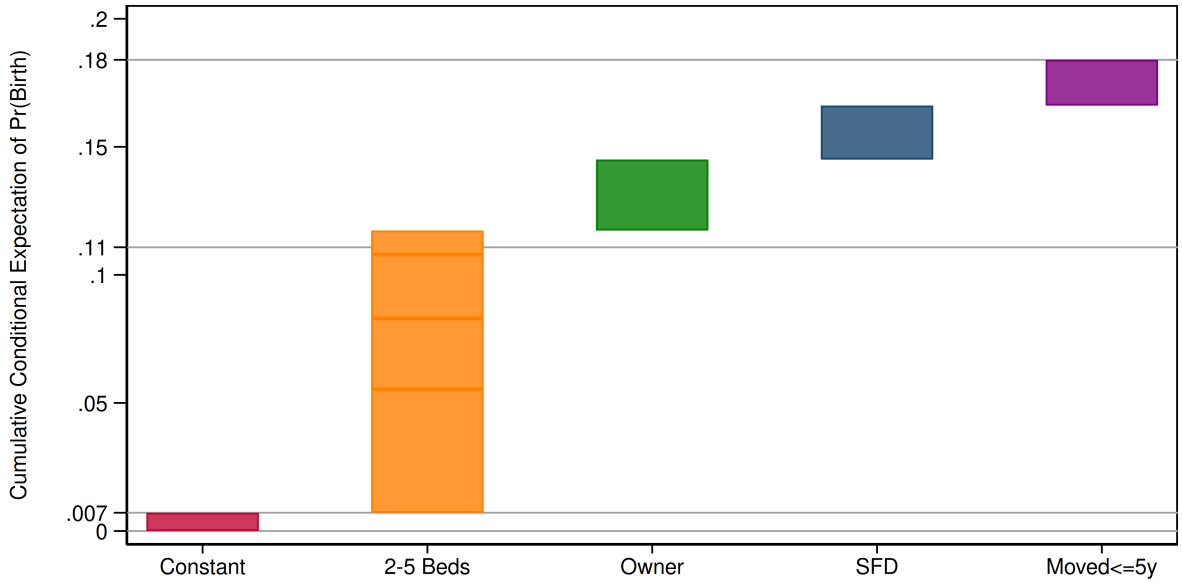


Figure 1: Results of a regression of having given birth in the last year on a vector of housing characteristics and city-year-age fixed effects: $\text{Birth}_i = \beta\ \text{Housing}_i + \text{FE}_{\text{city-year-age}(i)} + \epsilon_i$. Census and ACS microdata, filtered to women aged 26-33 who are household heads or partners of the household head. Results are arranged to show how the conditional expectation of the probability of having given birth increases when all housing variables are cumulatively switched to the family-friendly value: starting with the constant, moving from 1 bedroom up to 5 bedrooms, from renter to owner, from all other structures to single-family detached, and having moved to the unit in the last 5 years. The conditional expectation when all housing variables are set to the non-family-friendly values, at the mean, and when all housing variables are set to the family-friendly values, are depicted with grey horizontal lines.

The Relative Price of Family-Friendly Housing is Increasing. North American parenting norms dictate that each child have a separate bedroom, and so the normative demand for bedrooms is linear in fertility. However, in addition to the rise of housing costs in general, these larger sizes have been growing more expensive relative to one bedroom units. In Figure 2, I plot differences in year-bedroom fixed effects from a rent regression that controls for location. After a period of rent

compression over 1990-2000, 2 bedroom and especially 3+ bedroom units became substantially more expensive than 1 bedroom units.

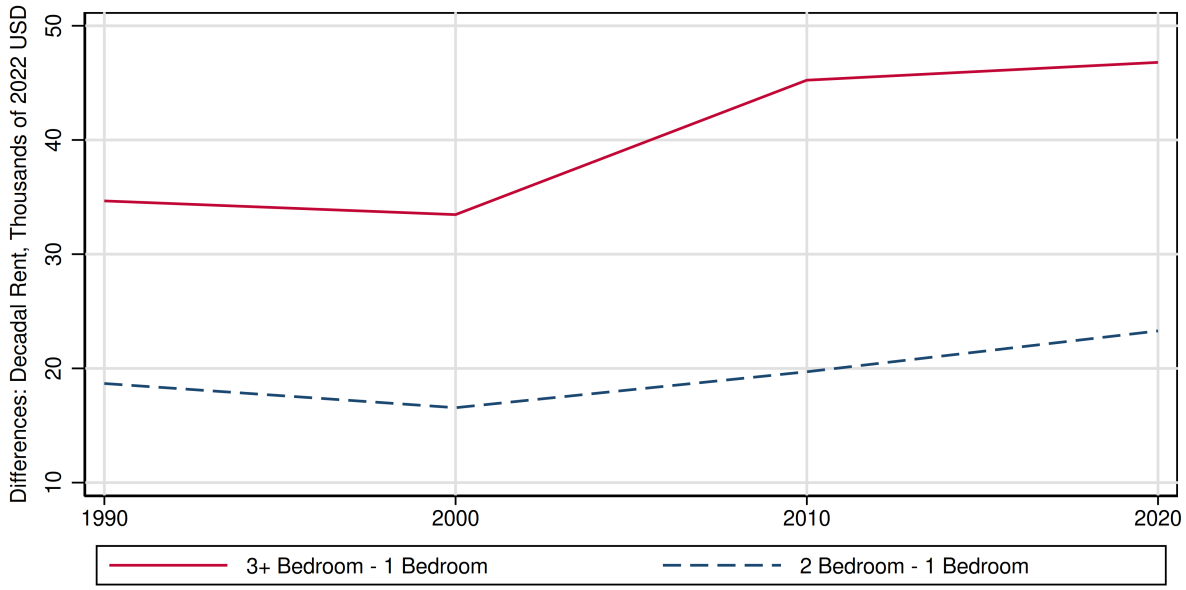


Figure 2: Results of a fixed effects regression: $\text{Rent}_{jbt} = \alpha_{bt} + \alpha_{jb} + \alpha_{jt} + \epsilon_{jbt}$, with differences in $\hat{\alpha}_{bt}$ plotted. Rent is measured in decadal rent (120 times monthly rent) in 2022 USD, the same measure used in structural estimation. Further details are in Section 2.1.

Household Size Drives Housing Demand. The size of a family household does not depend on the presence of minor children alone. Family households may be multi-generational, include adult children, include other adult family members or close friends, or may be headed by a parent whose relationship with the other parent has ended in divorce or separation. These other dimensions comprise additional sources of variation in the size of the household, and the size of the household is a primary determinant of the amount of living space demanded. In Figure 3, I plot by family size the CDF of bedroom count, which varies widely.

Household Formation is Declining, Driven by 20-29 Year Olds. Because very few children are conceived by parents who are themselves still living with their parents, a young adult's decision to live with their parents – to not form a new household of some kind – is tantamount to a decision not to have children. In Figure 4, I separate all adults (left panel) and 20-29 year olds (right panel) into (1) the head or the partner of the head of a family household, (2) a member of a non-family household (or in group quarters), and (3) a member of a family household that is not the head or partner of the head. The shifts in the overall adult population are large: a 6pp reduction in (1) split roughly equally between (2) and (3). In contrast, young adults have reduced (1) by 16pp, with about three quarters of the decrease flowing into (3) and only one quarter flowing into (2). Declining fertility is likely caused by declining household formation to some extent, and so it is important to

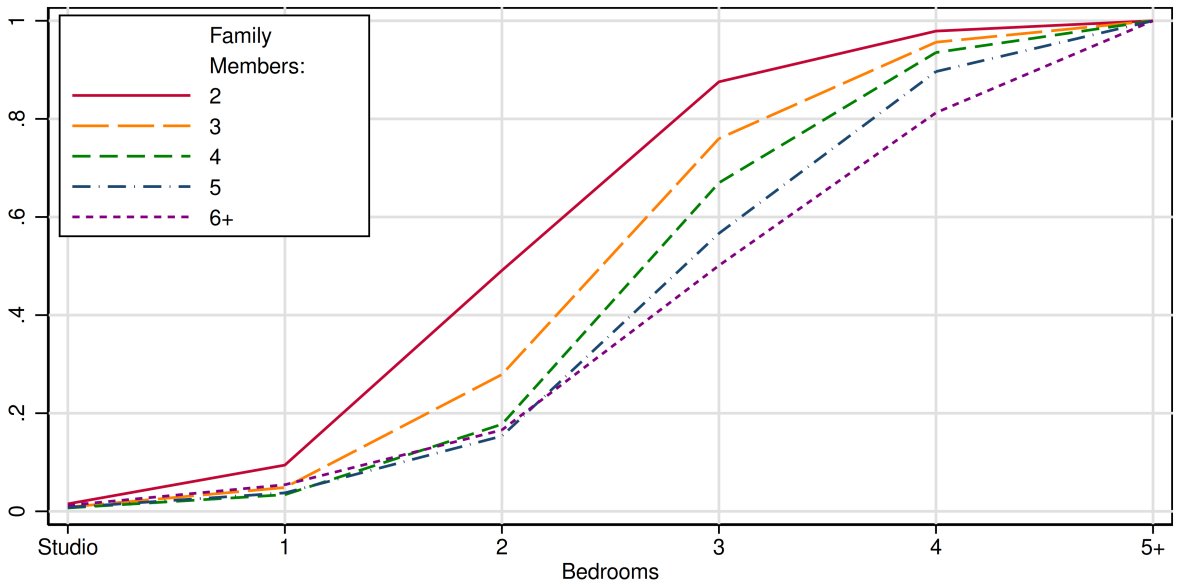


Figure 3: Cumulative distribution functions over bedrooms – from studio to 5 or more bedrooms – separately by family size. By definition, there is no such thing as a family household with only one member. Using pooled Census and ACS microdata, 1990-2022.

consider how housing costs may keep young adults living at home.

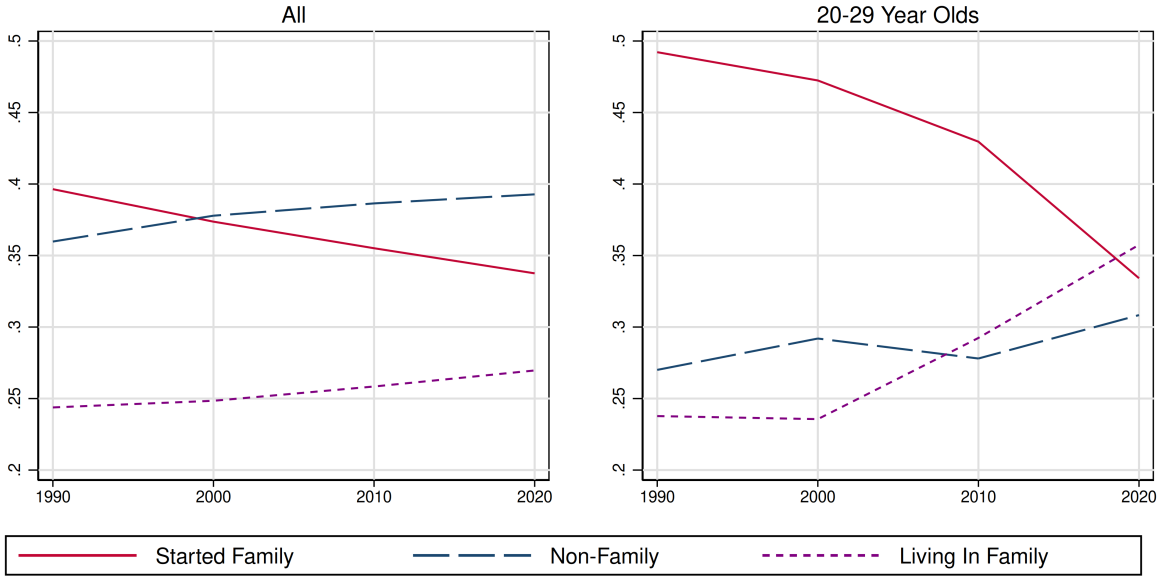


Figure 4: Shares of all adults (left) and 20-29 year olds in three categories: head or partner of the head of a family household, a member of a non-family household, or a member of a family household that is not the head or the partner of the head. Census and ACS microdata.

The Burden of Declining Household Formation Doesn't Just Fall on Parents. A key concept

in the model is the "extra" family member: a co-residing adult who is not a young adult child of the household head. However, these people will still serve to increase the size of the household and cause it to demand more housing. The share of people who are extra family members has increased by about 3pp from 1990 to 2020, and the share of family households that are hosting one has increased by about 9pp to nearly 30% in 2020. Accommodating this phenomenon is an important technical challenge of this paper.

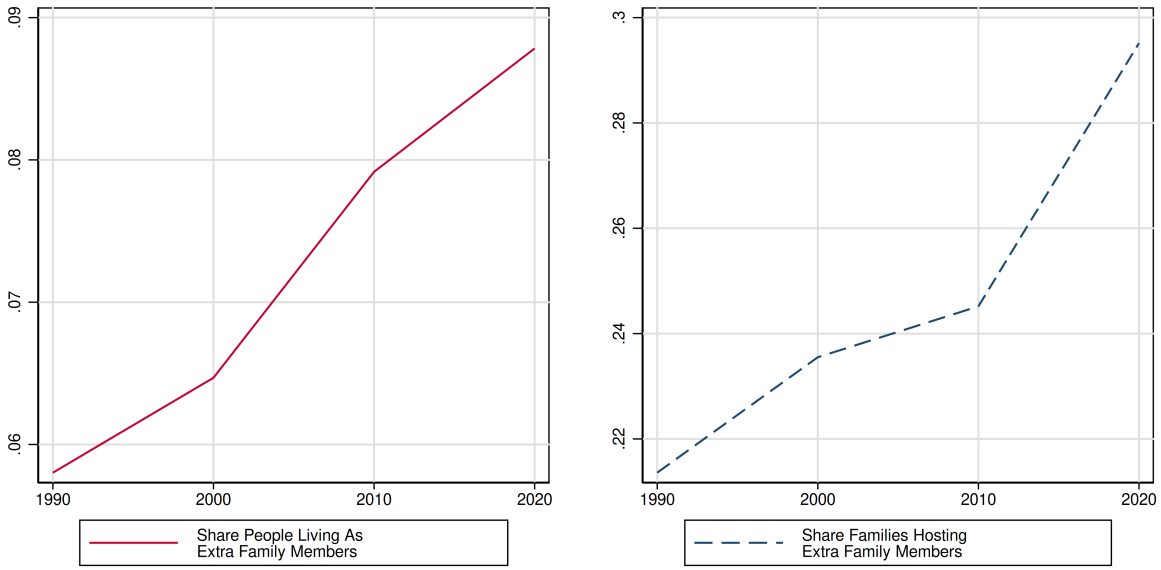


Figure 5: Extra family members are those who are not the head, partner of the head, or child of the head younger than 30. Left: share of population that is an extra. Right: share of family households that host an extra. Census and ACS microdata.

Non-Traditional Family Structures Are Increasing In Frequency. For centuries, Western life courses followed a simple sequence: birth to married parents in a two-adult household; departure in young adulthood; then marriage and children in another two-adult household. In recent decades this pattern has weakened under the "Second Demographic Transition," marked by "sustained subreplacement fertility, a multitude of living arrangements other than marriage, a disconnection between marriage and procreation, and no stationary population" (Lesthaeghe, 2014). Figure 6 traces the SDT's progress since 1990 in household structure. I define: (1) a traditional family as a household where the head's partner is present and the only other members are children up to age 19; (2) a non-family household as one where no two members are married or related; and (3) a non-traditional family as a family household where the head's partner is absent, a 20–29-year-old adult child co-resides, or an extra family member resides. The share living in a traditional family fell by 8pp since 1990, as 5pp shifted to non-traditional families and 3pp to non-family households. Traditional families, a plurality in the early 1990s, are now nearly 10pp below non-families and only

2pp above non-traditional families.

Absent the SDT, studying the effect of housing costs on fertility would be simpler because families would be more homogeneous. Instead, evidence already ties rising housing costs to delayed household formation among young adults and plausibly contributes to more multi-adult family arrangements. Ignoring the SDT would either restrict analysis to an ever-smaller, unrepresentative minority or miscode actual household sizes, distorting housing demand. I therefore specify a dynamic model of housing and fertility in which departures from the traditional structure are endogenous[3] and jointly determined with housing costs and fertility. Because unit size affects living-arrangement and fertility decisions differently, this structure is necessary to identify the fertility effects of policies targeting large versus small units – this paper's central aim.

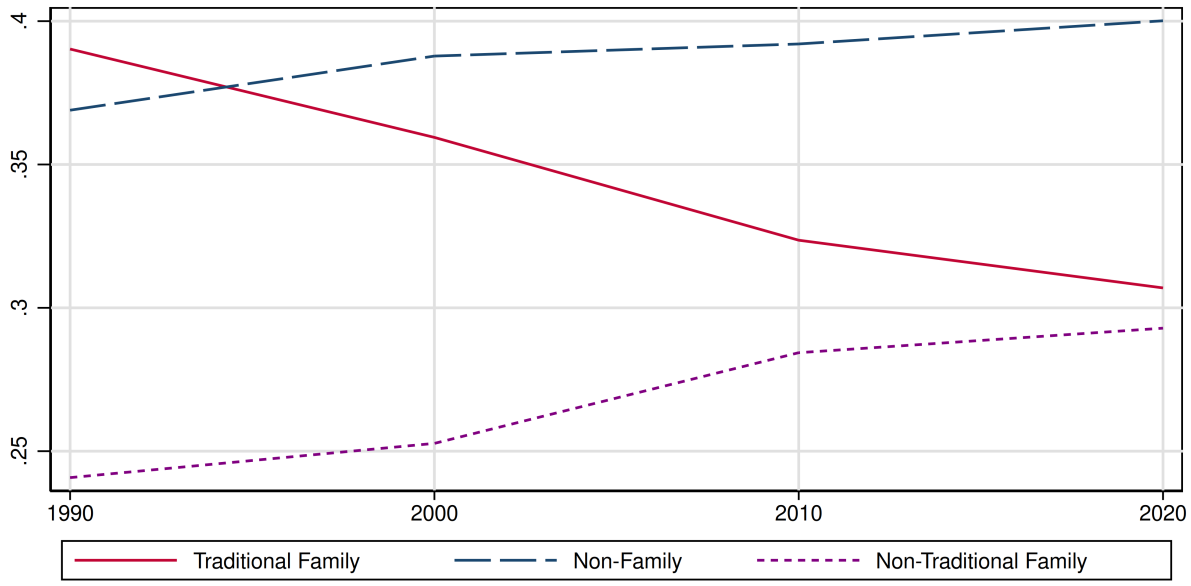


Figure 6: I use the Census Bureau definition of family households: one where at least two members are married or related by blood. I define "Traditional Families" as family households where the partner of the head is present and adult children and extra family members are not present, and "Non-Traditional Families" as family households where at least one of the preceding conditions does not hold. Census and ACS microdata.

3 Reduced-Form Analysis

In this section, I replicate the reduced-form negative relationship between housing costs and fertility, and show that fertility is more responsive to the cost of larger housing units. I use the same identifying variation as in structural estimation: shifts in the amenity value of *substitutable* locations

[3]Young-adult co-residence and extra family members are endogenous; partnership dissolution is exogenous and measured by to age-varying divorce rates over time.

that prompt equilibriating shifts in rents – the "donut IV" of Bayer *et al.* (2007). Finally, I discuss the reduced-form estimand and its limitations: even with experimental variation, the estimated causal effect of local housing costs on local fertility includes high-fertility families sorting into low-cost places and away from high-cost areas.

3.1 Regression Equation and OLS Results

I regress tract-level general fertility rates (GFR)[4] on the log of local median rents for 1 and 3+ bedroom units using the four-decade panel (1990, 2000, 2010, 2020) of consistent census tracts j described in Section 2.1. I present results using different sets of fixed effects – state, city, county, and tract, each with year fixed effects – and standard errors clustered by tract.

$$\text{GFR}_{jt} = \alpha_{\text{FE}(j),t} + \beta \ln \text{rent}^{b}_{jt} + \epsilon_{jt}, \quad b \in \{1, 3+\}. \tag{3.1}$$

The results of these regressions are presented in Table 1. Coefficients are stable across fixed-effects specification until tract fixed effects are added, which could result from the short panel (four decades). In each regression, the coefficient is negative and statistically significant, and the 3+ bedroom coefficient is between 1.6 and 3.6 times larger in magnitude than the 1 bedroom coefficient. Given an average GFR of 0.065, these correlations are relatively large: using Table 1 column (3), a doubling of 3+ bedroom rents is correlated with a decline in GFR that is equal to 25% of the mean.

Table 1: OLS regressions of the general fertility rate (GFR) on the rent of 1 and 3+ bedroom units, separately, with various fixed effects and standard errors clustered by tract.

	GFR			
Fixed Effects:	State, Year	City, Year	County, Year	Tract, Year
Ln(Med. Rent, 1Bed)	-0.00872***	-0.00892***	-0.00864***	-0.000866***
	(0.000154)	(0.000173)	(0.000183)	(0.000170)
Ln(Med. Rent, 3+Bed)	-0.0143***	-0.0160***	-0.0165***	-0.00310***
	(0.000228)	(0.000277)	(0.000303)	(0.000254)
N	177,926	177,926	177,926	177,926

[4]The GFR is the number of births per female of reproductive age, generally taken to be 15-44. I calculate it as the number of infants aged [0,1) in a tract divided by the number of women aged 15-44, scaled to match official GFR statistics from the CDC since the Census under-counts young children (discussed in greater detail in Section 2.1). In contrast, the total fertility rate (TFR) is a sum of age-varying fertility rates at a point in time, providing a measure of the lifetime fertility of a synthetic woman who lives her entire reproductive life at that point in time.

3.2 Identifying Variation and 2SLS Results

The endogeneity problem is severe and multidimensional. The objective is to understand the effect of rents on fertility, but this specification is afflicted by selection into treatment, heterogeneity over owners and renters, unobserved amenities that may disproportionately attract families or increase fertility, and a reverse-causal relationship as many large families increase demand for larger units. Although the reduced-form can address the latter two concerns, we do not observe separate small-area fertility rates by tenure, and even experimental variation cannot surmount the selection problem. Before introducing the structural model that will address selection and tenure heterogeneity, I introduce the identifying variation that I will use in structural estimation, and estimate a baseline relationship between housing costs and fertility that is causal up to those biases.

To generate identifying variation in rents, I turn to the donut instrument, the spatial version of the BLP (1995) instrument that was first introduced by Bayer *et al.* (2007). The donut IV is an average of amenities in a *donut*-shaped region around a *focal* tract. Neighborhoods sufficiently near are trimmed due to endogeneity concerns: people may have preferences over amenities in a region around the focal tract. Neighborhoods sufficiently distant are trimmed due to strength concerns: they are much less substitutable and provide little power in the first stage[5].

The donut instrument generates identifying variation by shifting the quality of substitutable (nearby) locations, which shifts willingness to pay for and the market-clearing rent of the focal tract. If the amenity value of substitutes increases, then rents in a focal tract must decline in order to restore equilibrium (and vice versa), provided that housing supply is not perfectly elastic. In structural estimation, the exclusion restriction is mean-independence of the donut instrument and unobserved characteristics of the focal tract, which is satisfied if the donut region is distant enough that any spatial correlation between the donut amenity and unobserved amenities of the focal tract has fallen to zero but near enough that tracts in the donut region are still substitutable with the focal tract. Although we cannot evaluate the spatial correlation between the instrument and the error term because the latter is unobservable, we can use a simple model and the facts that the median move distance is about 16-24km and about 75% of moves are within-metro to make an educated guess of about 20-30km for the median residential search radius (with details in Section A2). Too far beyond this distance it is implausible that the substitution mechanism is driving a first stage, and too far below we run the risk of violating the exclusion restriction.

I consider two amenities to use as the basis for a donut instrument: noise pollution from railways, and the share of housing units that are single-family detached homes. I also consider various distances over which to construct the instruments. It is important to emphasize once more that the identifying assumption is *not* that these amenities measured in the focal tract are valid instruments

[5]This is similar to Gandhi and Houde (2019), who construct stronger BLPL IVs using only products nearby in characteristic space. Physical proximity is a heuristic for characteristic proximity due to spatial autocorrelation.

for rent in the focal tract: then the exclusion restriction would be that the amenity has no effect on sorting utility outside of rents, which implies that it is not actually amenity and cannot form the basis for a donut instrument. A suitable amenity causes sorting to that location – hence the consideration of the share single-family, an amenity that was already in the sorting model. However, the share single-family is a proxy for many amenities which could be spatially correlated with the error term at longer distances, motivating a comparison with a more precisely-defined amenity: noise pollution. Noise pollution is difficult to consistently measure across an entire country at a fine spatial scale, but transportation networks can be modeled such that precise local estimates of transportation-based[6] noise pollution can be formed. Because railways are lines and not points, railway noise is more spatially autocorrelated and thus will have stronger spatial correlations with other things in the regression error term and require longer distances when calculating the donut instrument.

Given the regression specification and instrument definition, the 2SLS estimator is characterized by the following system of equations:

$$\begin{aligned} \text{GFR}_{jt} &= \alpha_{co(j),t} + \beta \ln(\text{rent}^b_{jt}) + \gamma x_{jt} + \epsilon_{jt} \\ \ln(\text{rent}^b_{jt}) &= \pi_z z^{\mathcal{D}}_{jt} + \pi_x x_{jt} + \pi_{co(j),t} + \eta_{jt} \\ z^{\mathcal{D}}_{jt} &= \frac{\sum_{k:d(j,k)\in\mathcal{D}} H_{kt} x_{kt}}{\sum_{k:d(j,k)\in\mathcal{D}} H_{kt}}. \\ \text{for} \quad b \in \{1, 3+\}, &\quad x_{jt} \in \{\text{single-family}_{jt}, \text{rail noise}_j\}, \quad \mathcal{D} \in \{[0,1), [1,2), ...\}. \end{aligned} \tag{3.2}$$

I show first stage regression results in Figure 7, where standard errors are clustered by census tract. For each combination of endogenous variable and instrument, the focal tract amenity (blue dashed line) has the right sign. Although the instrument has a coefficient of essentially zero when donut share single-family is used to instrument 3+ bedroom rents, all three other specifications have large regions where the coefficient has the right sign, is stable, and is strong according the Olea-Pflueger Effective F-statistic and critical values.

Second stage results are in Figure 8. The weak first stage when donut share single-family is used to instrument 3+ bedroom rents generates a very wide confidence interval (including zero), which I omit for visual clarity. In each case, the point estimate for the 3+ bedroom rent coefficient lies below the 1 bedroom rent coefficient, and the difference is statistically significant for distances below 25km when rail noise is used to construct the instrument. Coefficients are generally statistically distinct from zero, and the 1 bedroom coefficient is very similar regardless of which IV is used.

Because the individual donuts so far are all quite thin – a single kilometer in width, and a radius of many kilometers - I show results for consolidated donuts in Table 2. The rail noise instrument

[6]Road, rail, and aviation noise are all available, but road noise instruments are weak and aviation noise instruments fail sign diagnostics.

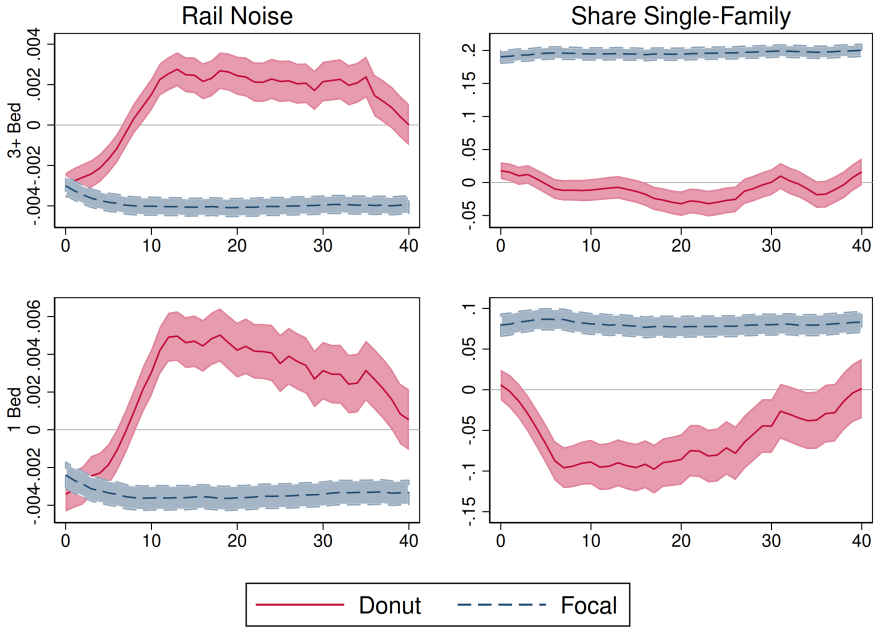


Figure 7: First stage regression coefficients (y-axis) for four models over a range of distances (x-axis). Left: rail noise is the IV; right: share single family is the IV; top: 3+ bedroom rents are explanatory variable; bottom: 1 bedroom rents are explanatory variable. The IV coefficient – amenity in the donut – is in solid red, while the coefficient on the amenity in the focal tract is in dashed blue.

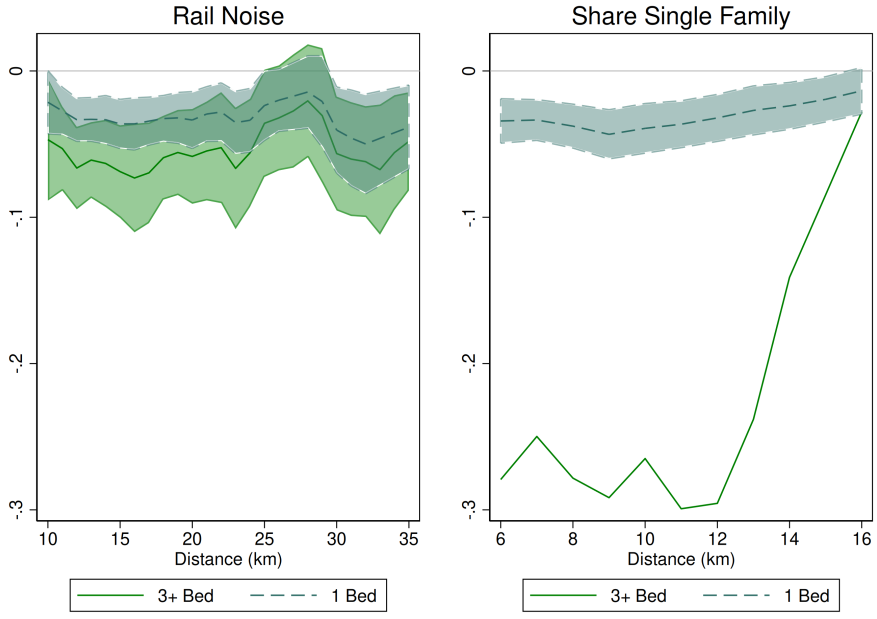


Figure 8: 2SLS coefficients (y-axis) for four models over a range of distances (x-axis). Left: rail noise is the IV; right: share single family is the IV. Very wide confidence intervals on one model omitted for visual clarity.

is strongest overall at a range of 11-15km, while the single-family instrument is strongest for 1 bedroom rents at a range of 6-9km. These distances align with the original findings of Bayer *et al.* (2007) that amenities (particularly, housing stock and land use) at a distance greater than 5km (3mi) tend to reduce focal house prices[7], and the idea that a wider radius is necessary when using rail noise since it is more spatially autocorrelated.

Baseline rail noise results are in columns 1 and 2; I add the share single-family as a control in columns 3 and 4, which slightly strengthens the first stage and reduces the magnitude of the coefficients. Baseline share single-family results are in columns 5 and 6; I add rail noise as a control in columns 7 and 8, which slightly strengthens the first stage and increases precision. I add tract fixed effects (infeasible for rail noise, which does not vary over time) in columns 9 and 10, which brings the 3+ bedroom rent coefficient into alignment with the rail noise results but causes the second stage coefficients to have the wrong sign for 1 bedroom rents (focal share single-family has the wrong sign in the first stage). In each case, the Anderson-Rubin test of significance of endogenous regressors rejects the null hypothesis, even in columns 5 and 7 when the instrument is weak.

The rail noise results suggest that, up to selection endogeneity, the semi-elasticity lies between -0.055 and -0.065 for large unit rents and -0.03 and -0.04 for small unit rents. The share single-family results, when the instrument is strong enough and first stage coefficients have the right sign (*i.e.*, ignoring columns 5, 7, and 10), indicate the same. These results are very large in magnitude relative to the general fertility rate average and interquartile range, 0.065 and 0.023, suggesting that sorting bias may be substantial.

3.3 Sorting Bias and SUTVA Violation

The reduced-form estimand is the causal effect of rents on fertility rates in a neighborhood, *without* holding constant the composition of the neighborhood, because large and small families sort differently over rent. If we are interested in understanding the aggregate effect of rents on fertility rates, we need a different estimand: the causal effect of rents on fertility, without a sorting bias.

Econometrically, specifying a reduced-form relationship between local fertility and local rents when (1) fertility and housing are chosen jointly, (2) there are latent preferences for fertility, and (3) fertility changes housing demand results in a violation of the Stable Unit Treatment Value Assumption. Under these assumptions, local fertility rates are a weighted average of the fertility rates of subpopulations with different latent preferences for fertility. The weights are determined by the distribution of latent fertility preferences in a location, which is determined by sorting and is a function of rents in other locations. When an outcome is a function of treatments in other units, SUTVA is violated: we are not estimating a treatment effect of rents on fertility rates that holds

[7] I discard the results that the single-family instrument is strongest for 3+ bedroom rents at a range of about 20km because it does not align with this pattern.

Table 2: Consolidated donuts based on first stage results with preferred specifications in columns 3 and 4. Baseline models in columns 1, 2, 5, and 6. I add the other amenity as a control in columns 3, 4, 7, and 8, and tract fixed effects (feasible only for the single-family instrument) in columns 9 and 10. Standard errors are clustered by tract.

	(1) GFR	(2) GFR	(3) GFR	(4) GFR	(5) GFR	(6) GFR	(7) GFR	(8) GFR	(9) GFR	(10) GFR
Log(Median 3+ Bed Rent)	-0.0646***		-0.0568***		-0.274		-0.242*		-0.0681***	
	(0.0149)		(0.0127)		(0.150)		(0.113)		(0.0119)	
Log(Median 1 Bed Rent)		-0.0327***		-0.0308***		-0.0406***		-0.0410***		0.0382***
		(0.00775)		(0.00747)		(0.00753)		(0.00736)		(0.00649)
Share Single-Family			0.0121***	0.00544***	0.0543	0.00598***	0.0471*	0.00610***	0.0128***	0.00564**
			(0.00246)	(0.000719)	(0.0293)	(0.000734)	(0.0213)	(0.000693)	(0.00286)	(0.00187)
Avg Rail Noise (dB)	-0.0000772	0.0000492	0.0000201	0.0000816**			-0.000524	0.0000498		
	(0.0000629)	(0.0000331)	(0.0000419)	(0.0000300)			(0.000335)	(0.0000307)		
Observations	186309	175035	186309	175035	186225	174984	186225	174984	186166	173945
IV	Rail, 11-15km	Rail, 11-15km	Rail, 11-15km	Rail, 11-15km	SF, 6-9km	SF, 6-9km	SF, 6-9km	SF, 6-9km	SF, 6-9km	SF, 6-9km
KP Stat	32.92	52.57	44.10	55.62	2.889	55.81	4.028	58.89	64.34	109.7
p-value: AR Test	0.00000388	0.0000123	0.00000530	0.0000247	1.59e-08	4.62e-09	3.64e-09	1.19e-09	4.33e-15	3.97e-13
Tract FEs									X	X

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

everything else constant. Rents affect the fertility rates of the subpopulations, but also the weight placed on each subpopulation since each group sorts to or away from that location at different rates due to their heterogeneous preferences over rent. The full argument is in Section A3.

I now provide empirical evidence that fertility and housing are chosen jointly in Table 3 (in addition to to Figure 1). First, I regress a birth indicator on a moving indicator, conditioning on fixed effects for year-age-tenure, current location, and location one year ago. This specification holds constant any location-specific factors – related to where she lives now, or where she did one year prior – that could affect fertility. I find that a woman who moves in a given year is 0.79pp more likely to also give birth in that year, which is an increase of 14.5% over the sample probability. Next, I regress a birth indicator on indicators for duration of residence in the current home, conditioning on fixed effects for year-age-tenure-city (past location is not observed when the residence duration exceeds one year). Here, the coefficient on moving in that year is about triple, and coefficients on having moved 1-2 and 2-5 years prior are slightly larger, indicating that moving in anticipation of having children is slightly more common than moving in the same year as having children.

I conclude that housing and fertility are chosen jointly, SUTVA is violated, and the reduced-form estimand includes sorting bias. I now turn to the structural model that is intended to remedy the bias.

Table 3: Reduced-form evidence that housing and fertility are chosen jointly. Moving is significantly correlated with giving birth in that year, and in the next several years. Using ACS microdata, 2005-2022. Robust standard errors.

	(1)		(2)	
	Birth in Last Year		Birth in Last Year	
Residence: [0, 1) Years	0.00786***	(0.000277)	0.0247***	(0.000295)
Residence: [1, 2) Years			0.0276***	(0.000353)
Residence: [2, 5) Years			0.0249***	(0.000230)
Residence: [5,10) Years			0.0127***	(0.000204)
Observations	11856932		11817377	
Mean LHS	0.0541		0.0541	
FE_1	Year-Age-Tenure		Year-Age-Tenure-City	
FE_2	PUMA, Current			
FE_3	PUMA, 1Yr Ago			

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

4 Model

The architecture of the structural model is motivated by the descriptive facts in Section 2, discussion of the reduced-form in Section 3, and now, the economics of housing and fertility. I begin by

discussing the results of extending Becker and Lewis's (1973) quantity-quality model to include housing. I then show how the structural model on a broad level accommodates these takeaways and provides a mechanism where high housing costs disincentivize fertility. Then, I discuss the details of the state, choice set, and state transition, which accommodate additional dynamic mechanisms of rising importance during the Second Demographic Transition. I then provide details about housing demand.

4.1 Quantity-Quality With Housing

After it was first sketched in Becker (1960), the quantity-quality tradeoff was formalized by Becker and Lewis (1973). The key element is that the cost of childrearing that enters the budget constraint is a product of the quantity of children and the pecuniary cost of the quality of children, where both quantity and quality also enter preferences. Then, if the income elasticity of child quality exceeds the income elasticity of child quantity, an increase in income will generate a reduction in fertility, even when the income elasticity of child quantity is positive.

To add housing to this framework, I define the housing choice variable as housing per child h, which enters preferences and creates an additional quantity-quality term in total housing expenditure nrh, quantity of children times cost of a unit of housing times housing per child. This term joins the standard quantity-quality term ni in the budget constraint, where per-child investment i is mapped to child quality q by an increasing function f. The parent's optimization problem is then

$$\max_{n,q,c,h} \mathcal{U}(n,q,c,h) \quad \text{subject to} \quad c + ni + nrh = y \quad \text{and} \quad q = f(i),\ f' > 0. \tag{4.1}$$

Under the assumptions that all goods are normal and have positive and declining marginal utility, we can easily see that the larger is the housing cost r, the more severe is the quantity-quality tradeoff related to housing, and the higher is the shadow price of fertility rh which reduces fertility via both the income and substitution effect.

The following results are proven in Section A4. Adding the assumptions that fertility is weakly separable from the other goods and that the elasticity of h and i with respect to n exceeds -1 (*i.e.*, that *total* housing and investment increases with the number of children), the optimal choice of all other goods decreases with n: less parental consumption, less per-child quality, less per-child housing. Then, the agent is pushed back onto a steeper part of the utility function, and thus has a stronger taste for the numeraire to spend on parental consumption and child quality, a stronger taste for a marginal unit of housing, and greater total housing consumption. These are the very patterns that the structural model is intended to exploit using heterogeneity in housing demand parameters.

4.2 Overview

This is a dynamic model of living arrangement (including fertility) and housing choice. Given a state and vector of iid preference shocks, the value is that of the best option in the choice set:

$$V(x, \epsilon^i) = \max_{g \in \mathcal{G}(x), h} \{v^x_{gh} + \epsilon^i_{gh}\} \quad \text{where} \quad (g, h, x) = (\text{living arrangement, housing, state}) \tag{4.2}$$

We assume that the flow of each option can be written as a sum of housing utility u and living arrangement utility α. We further assume that an agent's preferences for housing options depends on both their state entering the period and their choice. The value of each option is the sum of flow utility and the discounted expected future value of the option, which depends on the inclusive value of each future state and the probability of transitioning to each state. Notably, we assume that the housing choice does not affect these state transition probabilities:

$$v^x_{gh} = \underbrace{u_h^{\tau(g,x)} + \alpha^x_g}_{\text{flow utility}} + \underbrace{\delta \mathbb{E}\left[\sum_{x'} \underbrace{\pi(x' \mid g, x)}_{\text{state transition}} \times \underbrace{\ln\left(\sum_{h'} \sum_{g' \in \mathcal{G}(x')} \exp(v^{x'}_{g'h'}) \right)}_{\text{inclusive value}} \,\middle|\, \underbrace{W_{t(x)}}_{\text{info}_t} \right]}_{\text{continuation value} \equiv c^x_g}. \tag{4.3}$$

The assumptions that housing and fertility preferences are additively separable, and that the housing choice does not impact state transition, is essentially an assumption that moving costs on a decadal scale are negligible and that agents can frictionlessly re-optimize their housing choices in each period, which substantially improves tractability. The assumption that preferences over housing depend on living arrangement captures that the presence and quantity of children (and other dimensions of living arrangement) significantly affect how agents value housing options. We assume that preferences are linear over characteristics, with coefficients indexed by housing demand type τ, which is mapped into by the living arrangement state and choice:

$$u^\tau_h = X'_h \beta^\tau + \xi^\tau_h, \quad \text{where} \quad h = (\text{neighborhood, bedrooms}) \tag{4.4}$$

Heterogeneity in housing preferences is the means by which housing characteristics shape fertility choices. The quantity-quality model predicts that larger families will have a more negative rent coefficient, and a greater taste for additional bedrooms. A shift in rents for all units, or just the largest units, negatively impacts the value of being a large family the most.

However, as was shown in the descriptives section, the size of a family also depends on adult children and other family members, as well as whether both parents are present, and so all of these dimensions can also generate heterogeneity in housing preferences. If housing costs affect these

outcomes, there are additional channels that housing costs can affect the value of adding to the size of the family through having children. Omitting these channels would then provide an incomplete picture, and excluding these determinants of family size would require either using an estimation sample that is small and shrinking through the sample period or one in which a growing proportion of agents are inaccurately re-coded.

I thus specify a rich model of living arrangement choice and household formation. Adult children are agents that can remain with the family, form roommate households, or start families themselves. Agents can also choose to live in the household of some other family member. High housing costs disincentivize the formation of family or roommate households, causing agents to join other households, altering the housing preferences of the hosting decisionmakers. By increasing the likelihood that children remain with the family for multiple decades, high housing costs disincentivize having children, independently of the effect of childrearing on flow utility, and regardless of the inherent utility of having children. Likewise, starting a family increases the likelihood that other family members will unilaterally join the household, and high housing costs increase the probability of this outcome conditional on that choice, disproportionately hurting the value of having children through another channel. (I treat partnership dissolution as exogenous.) I describe the details of these dynamics in the next two subsections.

4.3 State and Choice Set

The complete state is

$$x = (\text{year, foreign/native-born, family, owner, age,}$$
$$\text{lost partner, kids aged 10-19, kids aged 20-29, extras})$$

Year, foreign/native-born, age, and lost partner are exogenous; everything else is endogenous. The unrestricted choice set is

$$\mathcal{G} = \{\text{group quarters, with parents, with other family, alone,}$$
$$\text{1 roommate, ..., } \bar{s}-1 \text{ roommates, have 0 kids, have 1 kid, ..., } \bar{n} \text{ kids, buy house}\}.$$

The first three options provide zero housing utility and can only be chosen in non-family, non-owner states. "Group quarters" means some form of institutional living (hospital, military barracks, college dorm, prison, *etc.*) and is included for completeness and to hold the tendency for people to live in such an environment – by choice or otherwise – constant in counterfactuals. "With parents" means choosing to remain with the household headed by the parents to which one was

born; it is available only in the second and third decade of life (first and second decade of being an agent with the ability to choose). "With other family" means living as an extra family member in some other family; availability of this option does not depend on age.

An agent can choose to live in a non-family household – that is, alone or with unrelated roommates – if the state is non-family.

Starting a family, taking a partner of the same age, and having one or more kids are all synonymous, and is available if the state is in the second through fifth decades of life. If the state is family or the choice is to start a family, an agent can also buy a house[8]. A family that has lost one of the partners cannot have additional kids.

4.4 State Transition

The year and age advance one decade per decade; foreign/native-born status is immutable.

The stock of children of each age evolves according to a binomial distribution that uses the previous-period, previous-age stock of children and the endogenous probability that adult children of that type chooses to live with parents. For the number of children aged 10-19,

$$Pr(n_{10-19} = k|x, n_{0-9,t-1}) = \left(\frac{n_{0-9,t-1}!}{k!(n_{0-9,t-1}-k)!}\right)\left(P_{\text{par}}^{10-19,t,\text{imm}(x)}\right)^{k} \times \left(1 - P_{\text{par}}^{10-19,t,\text{imm}(x)}\right)^{(n_{0-9,t-1})-k}, \quad (4.5)$$

and analogously for those aged 20-29.

The number of extras evolves according to a different binomial distribution. I assume that the total number of extras that live with households of a certain type is the product of the total number of people choosing to live as extras and a time-invariant share of extras going to each household type; this is the number of trials. I assume that every agent in a household type has an equal probability of hosting each extra so that the binomial distribution probability parameter is the reciprocal of the number of agents in that household type. The total number of extras is:

$$N_{\text{extras}}^{t} = \sum_{\tilde{x}} N^{\tilde{x}t} P_{\text{extra}}^{\tilde{x}t}. \quad (4.6)$$

The share of extras going to each "extras group" $w(\tilde{x})$ is:

$$s^{w(\tilde{x})} = \frac{\sum_t \sum_{\tilde{x}':w(\tilde{x}')=w(\tilde{x})} N^{\tilde{x}'t} e(x')}{\sum_t \sum_{\tilde{x}''} N^{\tilde{x}''t} e(x'')}. \quad (4.7)$$

[8]For simplicity, only families can buy a house. Otherwise, one roommate would be the owner and the other unrelated roommates would be paying rent – rare and difficult to model.

The number of extras going to households of a given type is the product:

$$N_{\text{extras}}^{w(\tilde{x}),t} = s^{w(\tilde{x})} N_{\text{extras}}^{t}. \tag{4.8}$$

Then the distribution uses this number of trials, and the reciprocal of the agent count in an extras group as the probability:

$$Pr(\text{extras} = e|x) = \left(\frac{N_{\text{extras}}^{w(\tilde{x}),t}!}{e! \left(N_{\text{extras}}^{w(\tilde{x}),t} - e \right)!} \right) \left(\frac{1}{N^{w(\tilde{x}),t}} \right)^{e} \left(1 - \frac{1}{N^{w(\tilde{x}),t}} \right)^{N_{\text{extras}}^{w(\tilde{x}),t} - e}. \tag{4.9}$$

There is an exogenous probability of a partnership dissolving that varies over time and age. If a partnership does dissolve, there is a 50% probability of being the partner that remains with the family, and a 50% probability of being the partner that reverts to a non-family state.

If all children leave the household and there are no extras then both partners revert to non-family. If they were owners, they revert to renters.

A non-family household may receive extras and become a family household for a decade.

4.5 Housing

Given the assumptions on the structure of v, a standard residential choice model can be derived, and I show this in Section A5. Because preferences for housing characteristics are heterogeneous in a way that is linked to fertility choice, changes in housing characteristics generate heterogeneous changes in values of fertility options and thus heterogeneous changes in fertility behavior. Particularly, larger families' stronger preferences over rent and size will combine with rising average and large-unit rents to generate decreases in fertility.

The utility of tract-bedroom $(j, b) = h$ housing option is linear in observables. I provide results for models with and without railway noise, but always include railway noise for exposition. Housing characteristics include the rent r, bedrooms b, single family share m, railway noise level x, a city-year-type fixed effect ϕ, and a demographic-varying unobserved utility ξ.

$$u_{jbt}^{\tau} = \beta_r^{\tau} r_{jbt} + \beta_b^{\tau} + \beta_m^{\tau} m_{jt} + \beta_x^{\tau} x_j + \phi_{c(j),t}^{\tau} + \xi_{jbt}^{\tau}. \tag{4.10}$$

Housing demand type has four dimensions, and an additively-separable structure is placed on the linear coefficients:

$$\tau = (\text{age, tenure, family/non-family, size}) = (a, e, f, s), \quad \beta^{\tau} = \beta + \beta^{a} + \beta^{e} + \beta^{f} + \beta^{s} \tag{4.11}$$

This structure assists with interpretability and is required for the later sufficient statistics results.

Notably, it is not placed on $\phi^{\tau}_{c(j)t}$, which is allowed to vary freely over (j, t, a, e, f, s). This fixed effect also proxies as a city-level shock to fertility utility, since it captures the value of being a large (or small) family in a specific city at a specific time.

5 Estimation

5.1 Housing Demand: Motivation and Discussion

Given that the necessary data for micro-BLP estimation is available – marginal distributions of household counts over neighborhoods and individual dimensions of demographic heterogeneity – using a standard estimator would be a natural approach, and would be a contribution to urban economics since it has not yet been done. However, the standard estimator constrains unobservable utility ξ_j to be homogeneous over demographics, and when the number of micro-moments is large this implies testable restrictions that hold with zero probability when the true data-generating process allows ξ^{τ}_j to vary over demographics. Instead, I use arguments pioneered by the Iterative Proportional Fitting (IPF) literature to show that these marginal distributions are sufficient statistics for the two-step Bayer *et al.* (2007) estimator when the residual varies over demographics[9].

As summarized by Conlon and Gortmaker (2023) there are many forms that micro-moments can take, some of which grow with the number of options J, and it is this growth with J that leads to model mis-specification. The form considered in this paper, and which are used in the sufficient statistics argument, is probability of choosing an option j conditional on one of the dimensions of demographics D taking value d, or $Pr(j|D = d)$. As the number of options J increases, the set of micro-moments increases by more than one (depending on the specification of demographics), while the number of parameters is constant and the number of error terms increases by only one. At a certain point, the parameters and errors – the data-generating process – will be unable to generate the observed data. The model will then be mis-specified, and the identification and consistency results that apply to correctly-specified models (Berry and Haile, 2014; Berry *et al.*, 2004b) will not apply. A formal proof of this argument is available in Section A6.

In addition to preventing mis-specification and providing certainty that foundational identification and consistency results apply, demographic-varying residuals are useful for several other reasons. First, abstracting from mis-specification, models with demographic-varying residuals explain the data more and demand less from the extreme-value shocks ϵ. The extreme-value shocks are assumed to be iid over individuals, but if a demographic-varying residual is non-zero, then we can conclude that there is unused information in the demographics that refutes that assumption. With demographic-

[9] *I.e.,* the marginal distributions provides equivalent estimates to those one obtains when using microdata or the tabulated joint distribution.

varying residuals, the extreme-value shocks are more likely to be truly idiosyncratic because we have conditioned on demographics to a greater extent. Second, without demographic-varying residuals, unobserved utility is a measure of purely vertical differentiation. Once demographic-varying residuals are added, unobserved utility may be primarily vertical, or different groups may have different rankings of the options – the data will determine which is true. Finally, the demographic-varying residual may be held constant in counterfactual analysis in order to better understand the differences between factual and counterfactual behavior for different groups.

I accommodate demographic-varying residuals – in the case where the joint distribution of choice and demographics is unobserved – by exploiting the sufficiency of the marginal distributions for the first-step maximum-likelihood estimator of Bayer *et al.* (2007)[10] under additive separability of preferences in coefficients (Equation 4.11). Heuristically, for any (non-error) term that enters the Poisson estimating equation, a count with the same indexes must be observed[11]. In Section A7 I prove this using a simple manipulation of the log-likelihood. Because the baseline utilities δ are identical, the second-step 2SLS estimator will return identical coefficients. Because the marginal distributions are sufficient, any joint distribution that aggregates into those marginal distributions will provide the same estimates as those of the true joint distribution when passed into an estimation routine that accepts rectangular datasets. I thus use IPF to combine the marginal distributions into

[10]This estimator is analogous to that of Berry *et al.* (1995) but with observed heterogeneity instead of unobserved heterogeneity in the sense that the heterogeneity and a baseline utility fixed effect are estimated in a first stage, and the baseline utility is decomposed to obtain the full distribution of preferences in a second step. In the case of BLP, we assume that the shocks on the coefficients are iid; in the case of BFM we assume that the demographic residual relative to the baseline residual $\xi_j^\tau - \xi_j$ is mean-independent of explanatory variables in order to consistently estimate heterogeneity in the first stage – in both estimators the instrument targets endogeneity in the shared unobserved utility ξ_j.

In this case, this assumption may be violated if neighborhoods that have amenities disproportionately valued by families (or which increase fertility) also have higher large-unit rents, or conversely if neighborhoods that have amenities disproportionately valued by singles (or which reduce fertility) also have higher small-unit rents. The former pattern would serve to reduce the extent of heterogeneity in rent-sensitivity and cause counterfactuals to be understated, while the latter would increase the extent of heterogeneity in rent-sensitivity and cause counterfactuals to be overstated. Because I include the share single-family in part to guard against this, and because it is *a prior* unclear which of the countervailing effects would dominate, it does not seem to be a major concern.

Regardless, one could avoid these concerns by estimating observed-heterogeneity logit models with 2SLS if the joint distribution of choices and demographics is observed and there are no finite sample zeros, or with the Poisson instrumental variables estimator of Mullahy (1997) if there are finite sample zeros. Although I have not proven it, it seems likely that the Mullahy (1997) estimator would remain consistent if the joint distribution of choices and demographics were not observed and an IPF-estimated joint distribution were used in its place.

[11]In a simplified context with options j and demographic heterogeneity over age a and sex s, the Poisson estimating equation is

$$N_j^{as} = \exp(\alpha^{as} + x_j(\beta + \beta^a + \beta^s) + \xi_j^{as})$$

where the fixed effect α^{as} captures both the logit denominator and the total number of agents of each type N^{as}, and is necessary for logit-Poisson equivalence (Figueirdo et al, 2003). By inspecting this equation, we can see what data is necessary: $(\alpha^{as}, x_j\beta, x_j\beta^a, x_j\beta^s) \implies (N^{as}, N_j, N_j^a, N_j^s)$, where notably N_j^a and N_j^s both aggregate into N_j and so only the former are required to be passed into estimation. This example shows the general pattern: we need to observe the marginal distributions of choice with individual dimensions of heterogeneity N_j^a and N_j^s, and the marginal distributions of heterogeneity but without choice N^{as}.

a single rectangular estimate of the joint distribution in a preliminary step. Because this step is equivalent to maximizing a likelihood, this estimated joint distribution is also sufficient for some linear estimators, which I prove in Section A8.

5.2 Housing Demand: Execution

As discussed above, I add a preliminary step to the two-step estimator of Bayer *et al.* (2007). I now discuss each step in turn: IPF estimation of the unobserved joint distribution, ML estimation of preference heterogeneity, and 2SLS estimation of baseline preferences.

5.2.1 IPF

I obtain an estimate of the tabulation of households H over tract j, number of bedrooms b, age of head a, tenure e, family/non-family indicator f, and size s for each of the four decades t in the panel by running the IPF algorithm[12] on a set of marginal distributions, where the first four marginal distributions are distributions of households over a subset of those dimensions and the fifth is the joint distribution at a higher level of geography: $p_t(j)$, the Public Use Microdata Area (PUMA) identifier from PUMA vintage t that consistent census tract j is mapped to. I run the algorithm in FORTRAN90, in parallel over PUMA-years, to a tolerance of 10^{-10}.

$$\hat{H}^{aefs}_{jbt} = \text{IPF}(H^{fs}_{jt}, H^{es}_{jt}, H^{af}_{jt}, H^{e}_{jbt}, H^{aefs}_{p_t(j),b,t}) \tag{5.1}$$

I additionally convert estimated households $\hat{H}$ to estimated decision-making agents N (Section A9.1) and remove some variation from rents so that the sufficiency result applies (Section A9.2).

5.2.2 MLE

For maximum-likelihood estimation of heterogeneity parameters, I exploit logit-Poisson equivalence[13] (Figueiredo *et al.*, 2003), use the total number of housing units H_{jt} as an exposure variable

[12] The algorithm consists of scaling cells in a guessed joint distribution as necessary to fit input marginals, looping over the set of marginals until convergence.

[13] As Figueiredo *et al.* (2003) showed, the logit and Poisson likelihoods have the same maximizers. This is convenient because it means that many logit estimators can be computed using more generic and performant routines that tolerate finite-sample zeros. It also raises the possibility of using the Poisson instrumental variable estimator of Mullahy (1997) when finite-sample zeros and endogeneity are present and joint estimation is desired.

(Section A9.3), and collect baseline utility into a single term δ_{jbt}:

$$\begin{aligned} N_{jbt}^{\tau} &= H_{jt} \exp\left(\beta_r^{\tau}\hat{r}_{jbt} + \beta_b^{\tau} + \beta_m^{\tau} m_{jt} + \beta_x^{\tau} x_j + \phi_{c(j),t}^{\tau} + \xi_{jbt}^{\tau}\right) \\ &= \exp\Big(\delta_{jbt} + (\beta_r^{\tau} - \beta_r) r_{jbt} + (\beta_b^{\tau} - \beta_b) + (\beta_m^{\tau} - \beta_m) m_{jt} \\ &\qquad + (\beta_x^{\tau} - \beta_x) x_j + (\phi_{c(j),t}^{\tau} - \phi_{c(j),t}) + (\xi_{jbt}^{\tau} - \xi_{jbt})\Big). \end{aligned} \tag{5.2}$$

The exposure variable is absorbed into baseline utility. Due to computational constraints, I randomly divide the data into four subsets[14] for estimation, obtain a consistent estimator for the linear heterogeneity parameters by averaging the estimate from either subset, and then partial out these parameters and regressors to obtain a consistent estimator for δ_{jbt} in a final Poisson likelihood maximization.

5.2.3 2SLS

Because the exposure variable was absorbed into mean utility, exposure can be corrected simply by subtracting $\ln H_{jt}$ from δ_{jbt}, and so the system of equations that defines the second step[15] is as follows:

$$\begin{aligned} \delta_{jbt} - \ln H_{jt} &= \beta_r r_{jbt} + \beta_b + \beta_m m_{jt} + \beta_x x_j + \phi_{c(j),t} + \xi_{jbt} \\ r_{jbt} &= \pi_b + \pi_m m_{jt} + \pi_x x_j + \pi_z z_{jt}^{\mathcal{D}} + \pi_{c(j),t} + \eta_{jbt} \\ z_{jt}^{\mathcal{D}} &= \frac{\sum_{k:d(j,k)\in\mathcal{D}} H_{kt} x_k}{\sum_{k:d(j,k)\in\mathcal{D}} H_{kt}}. \end{aligned} \tag{5.3}$$

Besides the outcome variable, there are other differences between this and the reduced-form specification. Here, the outcome variable is indexed by the number of bedrooms, allowing bedroom fixed effects and the rents of different size units to all enter the model. I allow rent to enter linearly – not in logs as in the reduced-form – so that preference heterogeneity captures the variation in local slopes that is implied by the quantity-quality model in Sections 4 and A4. I also use city fixed effects (CBSA, Core-Based Statistical Area) since the data used in IPF estimation is not sufficient for county fixed effects as in the reduced-form and PUMAs lack an inherent geographic concept.

[14] I thus conservatively adjust standard errors for the linear heterogeneity parameters by a factor of $\sqrt{4} = 2$ to accommodate the additional variance of the estimator.

[15] This system of equations describes the case where rail noise is the IV and x_j is in the model. When share single-family is the IV, rail noise is not in the model and $z_{jt}^{\mathcal{D}}$ is constructed using m_{jt}.

5.3 Living Arrangement and Fertility

The purpose of dynamic estimation is to estimate living arrangement utility α_g^x so that it may be held constant in decompositions and counterfactuals, thus isolating the effect of the housing market on fertility. There are four steps: kernel-smoothing of choice probabilities so that they are strictly between zero and one, inverting $\alpha_g^x + c_g^x$ from living arrangement choices given housing utility $u_h^{\tau(g,x)}$, forming estimates of the distribution of next-period values, and simulating c_g^x by taking draws from this distribution to finally back out α_g^x.

5.3.1 Kernel Smoothing

Due to the structure placed on preferences in Equation 4.3, dynamic estimation will primarily use the probability of of choosing a particular living arrangement option, aggregating over housing choices:

$$P_g^x = \sum_h P_{gh}^x. \tag{5.4}$$

The Hotz-Miller inversion will require probabilities that are strictly between zero and one. Due to a fine-grained state space, there are some states with a small number of agents and thus zero or one probabilities, and some states with no agents at all and undefined choice probabilities. I thus use a triangular kernel smoothing estimator with bandwidths 1.1, 2.1, 2.1, and 1.1[16] over decade, number of children 10-19, number of children 20-29, and extras respectively. I do not smooth over native/foreign-born status, age, tenure, family status, or lost partner status, because of the sharply different behavior of agents over these dimensions[17]. Hereafter, I refer to the smoothed probabilities as P_g^x to keep notation simple.

[16] I choose bandwidths ending in 0.1 because it limits the influence of other observations when smoothing for the focal observation is unnecessary, while giving other observations the same influence relative to each other when smoothing for the focal observation is necessary.

[17] Natives and foreign-born have very different fertility patterns, as do women of different ages. Because the focus of this paper is housing, I prefer to treat owners and renters as differently as possible. Families that are adding children are fundamentally different from agents making the choice to start a family, and if a family has lost one of its partners then it is incapable of having additional children. Meanwhile, I do allow smoothing over stocks of children and extras because a family having different levels of these stocks is less meaningful than some fundamental change of type. I smooth over decade because it seemed to be the least-bad marginal way to get all probabilities to be defined and strictly between zero and one.

5.3.2 Hotz-Miller Inversion

Due to the structure placed on preferences in Equation 4.3, we can factor fertility utility out of the fertility choice probability:

$$P_g^x = \sum_h P_{gx}^h = \frac{\exp(\alpha_g^x + c_g^x)\sum_h \exp(u_h^{\tau(g,x)})}{\sum_{h'}\sum_{g'\in\mathcal{G}(x)} \exp(\alpha_{g'}^x + c_{g'} + u_{h'}^{\tau(g',x)})}. \tag{5.5}$$

Then, we may log and difference with a reference choice for each state to isolate fertility utility:

$$\alpha_g^x + c_g^x = \ln\left(P_g^{zx} / \sum\nolimits_h \exp(u_h^{\tau(g,x)})\right) - \ln\left(P_{g_0(x)}^{zx} / \sum\nolimits_h \exp(u_h^{\tau(g_0(x),x)})\right). \tag{5.6}$$

Because there is no option that is available in every state, the values of different states are defined relative to a different option in each state. Specifically, the reference choice is living alone if the state is non-family, having zero children and remaining a renter if the state is renting family, and having zero children and remaining an owner if the state is owner family. There is one and only one reference choice in each state. I choose these reference choices because they are in some sense a default choice for each state, and the value of being a renting or owner family and having zero children is to some extent already contained in the value of transitioning to a renting or owner family by having at least one child and/or buying a house. In Section A11, I show that under an alternate normalization results are quantitatively different but overall robust.

5.3.3 Regression-Based Forecasting

Define the log of the living arrangement choice probability numerator as

$$v_g^{zx} = \alpha_g^{zx} + c_g^{zx} + \ln\left(\sum_h \exp(u_h^{\tau(z,x,g)})\right). \tag{5.7}$$

Letting $\tilde{x}$ denote the state without time t, assume that agents forecast future values as an autoregressive process with a shared coefficient, state-choice fixed effects, and state-choice time trends:

$$v_g^{\tilde{x},t+1} = \rho_0 v_g^{\tilde{x},t} + \rho_{g,1}^{\tilde{x}} + \rho_{g,2}^{\tilde{x}} t + e_g^{\tilde{x},t+1} \tag{5.8}$$

with $R^2 = 0.9891$. The distribution of errors will be drawn from to simulate the continuation value.

5.3.4 Simulating Continuation Value

We use simulation to estimate the expectation of non-linear functions. The continuation value for a particular draw d – where we have also substituted $v_{g'}^{x'}$ in pace of $v_{g'h'}^{x'}$ – can be written as

$$c_g^{xd} = \delta \sum_{x'} \pi(x'|g,x,d) \times \ln\left(\sum_{g' \in \mathcal{G}(x')} \exp(v_{g'}^{x'd})\right), \tag{5.9}$$

where the draw from the value function forecast error distribution impacts the future values in a straightforward way and the transition matrix in a more complex way through the endogenous choice probabilities of adult children and extras.

For the probability of partnership dissolution, I interpolate the divorce rate at different ages in different years to arrive at an age-varying decade-level divorce hazard rate[18].

For the probability of children remaining with the household when they are 10-19 (and analogously for 20-29), the draw affects the probability parameter in the binomial distribution:

$$Pr(n_{10-19} = k|x, n_{0-9,t-1}, d) = \left(\frac{n_{0-9,t-1}!}{k!(n_{0-9,t-1}-k)!}\right)\left(P_{\text{par}}^{10-19,t,\text{imm}(x),d}\right)^k \times \left(1 - P_{\text{par}}^{10-19,t,\text{imm}(x),d}\right)^{(n_{0-9,t-1})-k}. \tag{5.10}$$

The case of extras is thornier. To compute the probability of hosting an extra in the future, we must have an estimate of the total number of extras in the future, which depends on the extra choice probabilities (straightforward given the forecast and draw), and on the number of agents in each state in the future. It is probably possible, but very burdensome, to iterate over π and include external data on mortality and net migration to solve this using accounting properties:

$$N^{\tilde{x},t+1} = \sum_{\tilde{x}'} \sum_{g \in \mathcal{G}(x)} \pi(\tilde{x}, t+1|\tilde{z}', x', g, t) N^{\tilde{x}'t} P_g^{\tilde{x}'t} + \Delta_{\text{fertility}}^{\tilde{x},t+1} + \Delta_{\text{mortality}}^{\tilde{x},t+1} + \Delta_{\text{net migration}}^{\tilde{x},t+1}. \tag{5.11}$$

Instead, I take agents to form a deterministic forecast of the period-ahead distribution of agents over states based on a regression similar to that used to forecast values:

$$N^{\tilde{x},t+1} = \psi_0 N^{\tilde{x}t} + \psi_1^{\tilde{x}} + \psi_2^{\tilde{x}} t + e^{\tilde{x},t+1}. \tag{5.12}$$

[18] 1990, 2010, 2012, 2015, 2017, 2019, and 2021 compiled by the National Center for Family and Marriage Research at Bowling Green University from CDC and ACS data (NCFMR, 2012; NCFMR, 2014; NCFMR, 2017; NCFMR, 2019; NCFMR, 2021; NCFMR, 2023).

Agents then approximate

$$N^{\tilde{x},t+1} \approx \tilde{N}^{\tilde{x},t+1} = \hat{\psi}_0 N^{\tilde{x}t} + \hat{\psi}_1^{\tilde{x}} + \hat{\psi}_2^{\tilde{x}} t, \tag{5.13}$$

which I consider to be an acceptable approximation since $R^2 = 0.9985$. Simulation draws for value function errors are used in the period-ahead choice probabilities:

$$\tilde{N}_{\text{extras}}^{\tilde{x},t+1,d} = s^{w(\tilde{x})} \sum_{\tilde{x}'} \tilde{N}^{\tilde{x}',t+1} P_{\text{extra}}^{\tilde{x},t+1,d}. \tag{5.14}$$

And finally, the probability distribution for extras conditional on an agent's state and a vector of simulation draws:

$$Pr(\text{extras} = e|\tilde{x}, t+1, d) = \\ \Big(\frac{\tilde{N}_{\text{extras}}^{w(\tilde{x}),t+1,d}!}{e!(\tilde{N}_{\text{extras}}^{w(\tilde{x}),t+1,d} - e)!}\Big)\Big(\frac{1}{\tilde{N}^{w(\tilde{x}),t+1}}\Big)^e\Big(1 - \frac{1}{\tilde{N}^{w(\tilde{x}),t+1}}\Big)^{\tilde{N}_{\text{extras}}^{w(\tilde{x}),t+1,d} - e}. \tag{5.15}$$

With $Pr(n_{10-19} = k|\tilde{x}, t+1, d)$, $Pr(n_{20-29} = k|\tilde{x}, t+1, d)$, and $Pr(\text{extras} = e|\tilde{x}, t+1, d)$ known, we also know $\pi(x' \mid x, g, d)$, and we can finally estimate α as

$$\hat{\alpha}_g^x = \alpha_g^x + c_g^x - \delta \sum_d \frac{1}{D} \sum_{x'} \pi(x' \mid x, g, d) \times \ln\left(\sum_{g' \in \mathcal{G}(x')} \exp\left(\hat{v}_{g'}^{x'd}\right)\right). \tag{5.16}$$

6 Equilibrium

In the decompositions, I manipulate rents directly, holding constant the intrinsic fertility utility and recomputing the continuation value to solve a partial equilibrium. In the counterfactuals, I add and shift a housing supply equation, also solving for rents in general equilibrium.

6.1 Partial Equilibrium and Decomposition

A central secular trend in the last few decades in the US and other developed economies is rising housing costs. Since housing is often the largest part of a household's budget, this trend has far-reaching consequences. Effective housing policy, which has been lacking, would restrain the growth of housing costs by ensuring that housing supply is elastic so that new units can be built at low cost wherever rising demand pushes rents and prices up. Given the evidence about the relationship between housing costs and fertility presented thus far, the question of what fertility rates would be in a world where housing costs are not secularly increasing is of first-order importance.

The reduced-form literature attempts to answer this question using quasi-experimental variation in endogenous housing costs, but sorting limits the relevance of these estimands. Therefore, the first set of exercises in this paper are decompositions that target rents directly and can be considered partial equilibrium exercises. I first set rents equal to their values in 1990 in order to understand what the total effect of rising housing costs is. Because agents respond in both their housing and living arrangement decisions, the quantity of housing units demanded changes over size and location. Since quantities change while rents stay flat, this decomposition can also be thought of as a counterfactual where housing supply is perfectly elastic, and an upper bound on the extent to which housing policy can increase fertility. Then, I let average rents rise while relative rents are fixed and vice versa to understand whether rising average or relative rents are more important.

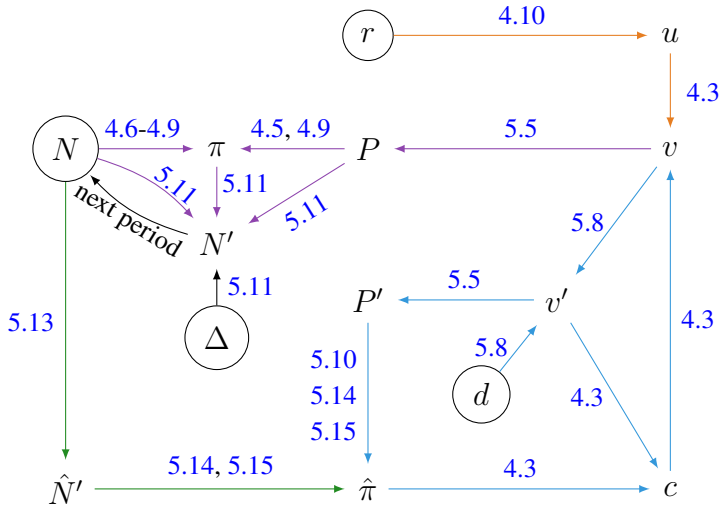


Figure 9: Visual partial equilibrium definition and solution algorithm. Initial shift in orange, inner loop in blue, the two parts of the outer loop in purple and green, passing of the actual distribution of agents over states forward in time in black. Circled nodes are inputs; N is circled despite having an arrow pointing to it because for each decade it is an input that is an output of the prior decade.

Figure 9 is a visual depiction of the partial equilibrium definition, showing which variables affect which and referring to the specific equations that define each relationship (a more formal partial equilibrium definition is available in Section A10). To solve this partial equilibrium fixed point problem, I first shift rents r which affects housing utility u and value functions v, shown in orange. In an inner loop depicted in blue, I update future values v' based on the forecast regression parameters and draws from the error distribution. The future values directly affect the continuation value c, but also affect future choices P' in the form of the decisions of adult children and extras. These decisions affect the *predicted* state transition matrix of other agents $\hat{\pi}$, which also affects continuation value c, which feeds back into current values v. Before each inner loop initiates, an outer loop completes an iteration of two distinct calculations. In purple, current choices P and state distribution N are

used to form the *actual* state transition matrix π in order to update the *actual* distribution of agents over states N' in every period with an accounting relationship (including the exogenous effect of mortality and net migration Δ, calculated as the residual between the model-predicted and actual distribution of agents over states from a baseline run of the decomposition using factual inputs), which then becomes the actual distribution of agents over states in the next period N (in black). Once the actual distribution of agents over states is known, it is used (in green) to form agents' *predictions* of the future distribution of agents over states $\hat{N}'$ via a reduced-form relationship, which is then combined with the *predicted* future choice probabilities P' in every inner loop iteration to obtain the distribution of extras in the expected state transition matrix $\hat{\pi}$.

After solving the decomposition, I compute children born and cohort and total fertility rates, as well as household formation and housing market outcomes.

6.2 General Equilibrium and Counterfactuals

Because housing costs are endogenous, they cannot be directly controlled by policy, and so to understand the effects of policy going forward we need to consider general equilibrium counterfactuals. Much of the focus of housing policy reformers in recent years has been on increasing building of any kind, agnostic to any heterogeneity in housing units. However, regulations and market incentives combine to make smaller units lower cost and more profitable and thus the primary result of any YIMBY policy, which – as has been discussed elsewhere in this paper – has an ambiguous effect on fertility[19]. Thus other policy experts have advocated a more specific focus on larger units in order to arrest the increase in the cost of starting a family. This policy dilemma – focus on building small and low-cost units, or larger and more expensive units – is the focus of the policy counterfactuals.

The policies considered are shifts in housing supply curves for 1 bedroom and 3+ bedroom units. These cost reductions can be thought of as developer subsidies, or reductions in the regulatory tax of building. To keep the policies as comparable as possible on a cost basis, I consider supply shifts that would result from equal expenditures – 5% of total rent expenditures in each year, treating owners as renters – on subsidies at initial rents and quantities.

The housing supply equation is:

$$\ln r_{jbt} = \gamma_{jb} + \gamma_{jt} + \gamma_{bt} + \gamma_j^{-1} \ln H_{jbt} + \epsilon_{jbt}. \tag{6.1}$$

I take the elasticities from Baum-Snow and Han (2024), and I estimate the heterogeneous intercepts in stages[20]. Fertility decisions feed back into housing quantities via the following four equations.

[19] More affordable small units may increase fertility by drawing singles away from large units or by reducing the long-run cost of fertility as adult children are prompted to move out, but may also increase the value of living in a small unit as a single relative to starting a family.

[20] If I attempt to estimate this equation jointly (including as a regressor or partialling out $\gamma_j^{-1} \ln H_{jbt}$) then larger units

Reading from right to left, first, I aggregate over dynamic states and choices to obtain the total number of agents for each housing demand type, then I find the tract-bedroom distribution of each type using the housing choice probabilities, then I convert to households using the rule of thumb discussed in Section A9.1, and finally I aggregate over housing demand type to find total housing demand, which enters the supply equation.

$$H_{jbt} = \sum_{\tau} H^{\tau}_{jbt}, \quad H^{\tau}_{jbt} = \begin{cases} N^{\tau}_{jbt}/s(\tau) & \text{if} \quad f(\tau) = 0 \\ N^{\tau}_{jbt}/2 & \text{if} \quad f(\tau) = 1, \end{cases} \quad N^{\tau}_{jbt} = P^{\tau}_{jbt} N^{\tau}, \quad N^{\tau} = \sum_{(g,x):\tau(g,x)=\tau} P^{x}_{g} N^{x} \tag{6.2}$$

As shown in Figure 10, the general equilibrium definition (a more formal equilibrium definition is available in Section A10) adds these conditions to the partial equilibrium definition, and the counterfactual solution algorithm adds a market-clearing loop to the decomposition solution algorithm. Now, after housing supply parameters γ shift rents r, housing utility u, and value functions v, the outer loop computes the actual and expected future distribution of agents over states N' and $\hat{N}'$, and the inner loop finds the inner fixed point v, I use the distribution of agents over states and types N and living arrangement and housing choice probabilities P_g and P_h to find the distribution of households over housing options H, implying a new rent r. This rent shifts utility, housing probabilities, and the number of households, generating a new rent, *etc.* until convergence and the next outer loop iteration.

7 Results

I first present housing demand results which align with the predictions of the Becker quantity-quality model extended to include housing, and which govern the relationship between housing and fertility choices. I then illustrate the causal effect of housing costs on fertility and household formation via a partial equilibrium decomposition. Finally, I show how alternate policies affect fertility and household formation via two general equilibrium counterfactuals.

7.1 Housing Demand

Matching theoretical predictions, structural estimates show that larger families are more sensitive to rents and have a stronger taste for large units. This pattern of heterogeneity ensures that rising rents – in general and for the largest units – reduce fertility, and is estimated in a first step with

have a lower intercept despite having higher rents, which is probably related in some way to the fact that the elasticity is constrained to be identical over size. Instead, I estimate this equation in stages. I first include $\gamma_j^{-1} \ln H_{jbt}$ as a regressor to estimate the fixed effects, and I retain the most important one γ_{bt}, which will be targeted by counterfactuals. I then partial out $\gamma_j^{-1} \ln H_{jbt}$ as if it had a coefficient of one and γ_{bt} before estimating the other fixed effects γ_{jb} and γ_{jt}. This way allows the tendency for larger units to have higher rents to pass through into the intercepts, while still allowing the other fixed effects and elasticity term to have some influence.

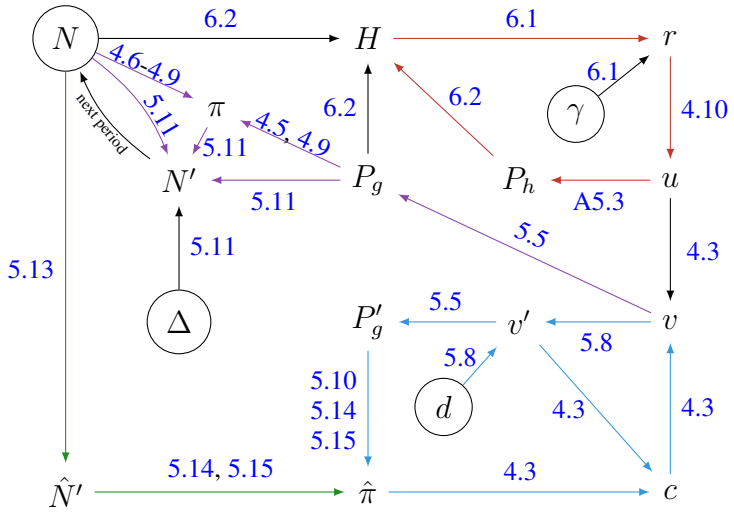


Figure 10: Visual general equilibrium definition and solution algorithm. Initial housing supply shift in black ($\gamma \to r$), then $r \to u \to v$ and outer loop initiates in purple. The actual distribution of agents over states N is obtained in black, then next part of outer loop in green. Inner loop in blue iterates until convergence, then rent loop in red initiates and upon completion another outer loop begins in purple, *etc*.

maximum likelihood. The absolute values of the rent coefficients determine the steepness of the demand curves and the quantitative magnitudes of these effects, and is estimated in a second step with two-stage least squares.

A type's rent coefficient is the sum of a baseline coefficient and a distinct heterogeneity term for every level of every dimension:

$$\beta_r^\tau = \overbrace{\beta_r}^{\text{baseline}} + \overbrace{\beta_r^a}^{\text{age}} + \overbrace{\beta_r^e}^{\text{tenure}} + \overbrace{\beta_r^f}^{\text{family}} + \overbrace{\beta_r^s}^{\text{size}}, \tag{7.1}$$

with one level of every dimension normalized to have a zero heterogeneity term. Preferences for bedrooms, share single family, and rail noise follow the same pattern. [21] The heterogeneity terms are estimated in the first step, and the baseline coefficient that shifts the entire distribution of preferences is estimated in a second step.

I begin by presenting first stage results for the 2SLS model that comprises the second step. In Figure 11, the top two panels depict Olea-Pflueger Effective F-statistics, the bottom two panels depict first stage coefficients, the left two panels use the rail noise instrument, and the right two panels use the share single-family instrument, all over a range of distances used to construct the donut instrument. The first stage is generally stronger than in the reduced-form analysis, but since errors are clustered by tract this is not merely a result of using all three bedroom sizes and thereby

[21] This structure aids interpretability, but is also required to some extent when using micro-moments for identification, depending on which micro-moments are used. In contrast, the city-type-year fixed effect ϕ_{ct}^τ is allowed to vary freely over τ, because counts by city-year-type are used.

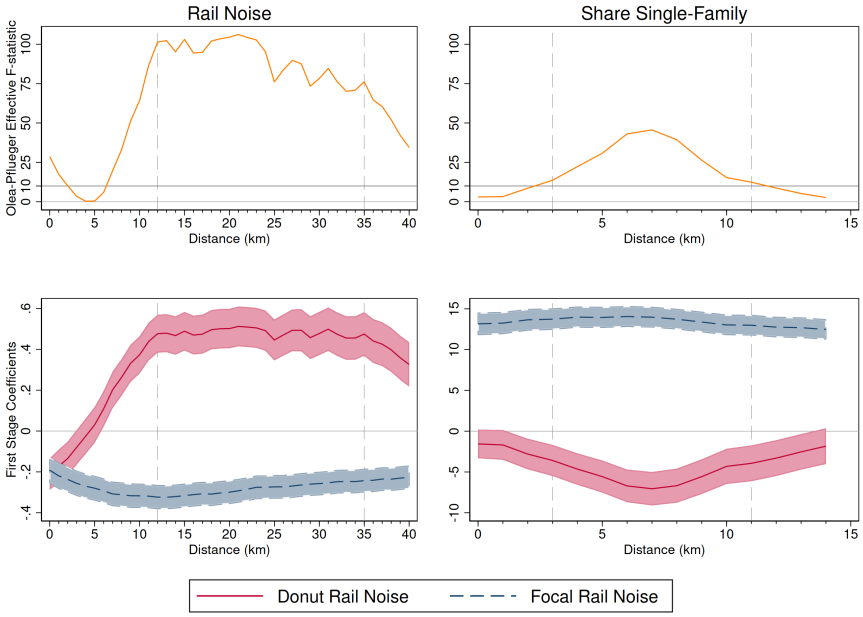


Figure 11: First stage regression results in housing demand estimation. The IV is based on rail noise at left and share single-family at right. Top panels are Olea-Pflueger Effective F-statistics, and bottom are coefficients. Standard errors are clustered by tract.

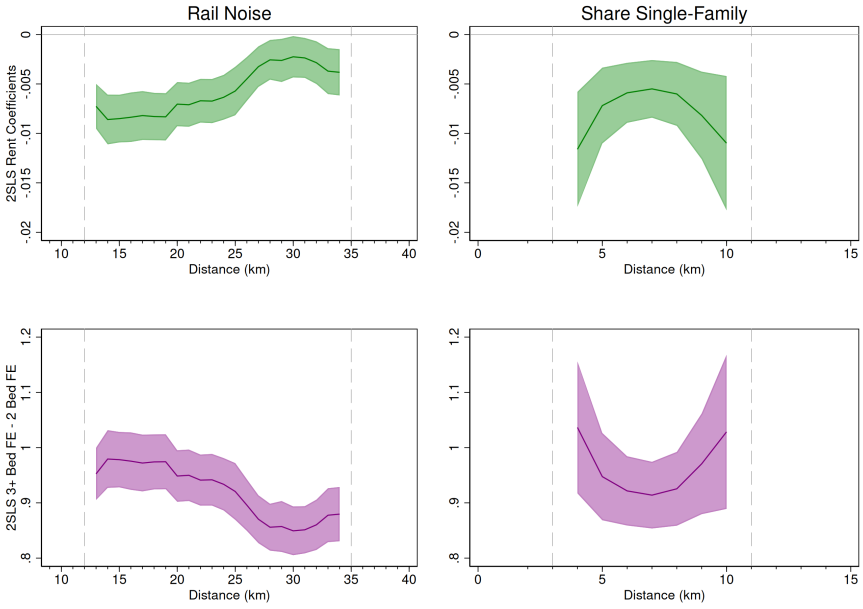


Figure 12: Second stage regression results in housing demand estimation. The IV is based on rail noise at left and share single-family at right. Top panels show the rent coefficient, and bottom panels show the difference between the 3+ and 2 bedroom fixed effects. Standard errors are clustered by tract.

increasing the sample size. For either instrument there is a range of distances where the first stage coefficients are stable and the Olea-Pflueger Effective F-statistics exceed 10. As in the reduced-form analysis, this range is closer to a distance of zero for the share single-family instrument.

Second stage results are presented in Figure 12. Again, the left panels show results using the rail noise instrument, and the right panels show results using the share single-family instrument. The top panels show the rent coefficient, while the bottom panels show the difference between the 3+ bedroom fixed effect and the 2 bedroom fixed effect. Results are truncated to fall within the region where the instrument is strong and has the right sign. There are two notable patterns in this figure. First, both instruments provide estimates that are quantitatively similar, despite the differences in how they were constructed, the range where they are strong, and the patterns of variation over distance. Second, the difference between the 3+ and 2 bedroom fixed effects is a mirror image of the rent coefficient. Upon reflection, this is not surprising: since larger units have higher rents, choices must either be rationalized by either a taste for the good or a distaste for its cost.

Next, I aggregate donut regions where the instruments have particularly strong and stable first stages, and present OLS, first stage, and 2SLS results in Table 4. For the rail noise instrument, I choose 12-24km, where the KP statistic hovers around 100 and the the donut coefficient is close to constant. For the share single-family instrument, I choose 5-9km, which brackets the point where the donut is strongest and where the KP statistic is 25 or higher. I also include the single kilometer donut that generates the most negative rent coefficient for each instrument.

There are two OLS columns, because the two left hand side variables δ_{jbt} come from two maximum-likelihood estimates, one that includes focal tract rail noise and one that does not – however the difference is apparently very small. Turning to the first stage estimates, using wider donuts tends to increase the coefficient and the standard error for a net increase of the KP stat. Wider donuts also generate second stage coefficients that are smaller in magnitude. In every 2SLS specification, the Anderson-Rubin F-statistic verifies the statistical significance of the rent coefficient.

In Figure 13 I plot histograms of the parameters of interest (weighted by the frequency of agents of different housing demand types) given baseline preferences as in Column 5 (rail instrument, 14km) in Table 4, the specification that generates the most negative rent coefficient. All types have downward-sloping demand and value 2 bedroom units more than 1 bedroom units. However, there is a small mass that values 3+ bedroom units less than 2 bedroom units. As shown in Figure 12, there is an inverse relationship between the difference between the 3+ and 2 bedroom fixed effect and the rent coefficient. Even the most negative rent coefficient is insufficient for all agents to value additional space, which we know must be true because it is an unambiguous good. I conclude that the true rent coefficient is even larger, but take the estimate from Table 4 Column 5 as the preferred estimate for the decompositions and counterfactuals since it is closest to the truth. This bias reduces the magnitude of the decomposition and counterfactuals, making them more conservative.

Table 4: OLS, first stage, and second stage results using both preferred single-kilometer donuts and consolidated donuts for the two IVs. Standard errors clustered by census tract. Preferred specification: Column 5.

	(1) OLS	(2) First	(3) 2SLS	(4) First	(5) 2SLS	(6) OLS	(7) First	(8) 2SLS	(9) First	(10) 2SLS
Rent	-0.000139**		-0.00621***		-0.00899***	-0.000160***		-0.00723***		-0.00543***
	(0.0000467)		(0.000800)		(0.00136)	(0.0000466)		(0.00152)		(0.00138)
2 Bedrooms	1.143***	19.10***	1.259***	19.10***	1.313***	1.144***	19.10***	1.279***	19.10***	1.244***
	(0.00411)	(0.0344)	(0.0158)	(0.0344)	(0.0262)	(0.00411)	(0.0344)	(0.0291)	(0.0344)	(0.0265)
3+ Bedrooms	1.948***	39.65***	2.189***	39.65***	2.300***	1.949***	39.65***	2.230***	39.65***	2.158***
	(0.00680)	(0.0481)	(0.0324)	(0.0481)	(0.0542)	(0.00679)	(0.0481)	(0.0604)	(0.0481)	(0.0549)
Share Single Family	-0.722***	12.46***	-0.650***	12.23***	-0.617***	-0.726***	14.09***	-0.639***	13.93***	-0.661***
	(0.00703)	(0.613)	(0.0120)	(0.613)	(0.0182)	(0.00702)	(0.624)	(0.0209)	(0.623)	(0.0188)
Rail Noise, Focal	0.00272***	-0.338***	0.00103**	-0.322***	0.000258					
	(0.000259)	(0.0290)	(0.000367)	(0.0291)	(0.000508)					
Rail Noise, 12-24km		0.836***								
		(0.0608)								
Rail Noise, 14km				0.461***						
				(0.0488)						
Share Single Family, 5-9km							-8.342***			
							(1.157)			
Share Single Family, 7km									-7.726***	
									(1.062)	
N	424659	424659	424659	424659	424659	424659	424659	424659	424659	424659
Instrument			Rail 12-24km		Rail 14km			SF 5-9km		SF 7km
Olea-Pflueger F-Statistic			189.3		89.21			52.02		52.94
Anderson-Rubin F-Statistic			81.48		78.84			34.84		19.66

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

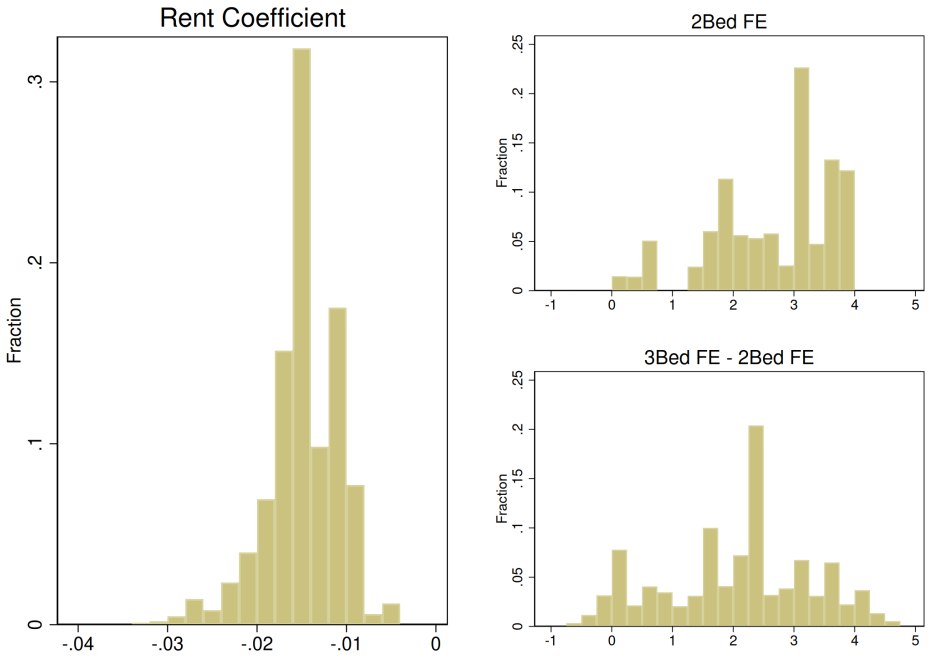


Figure 13: The distribution of preferences over type given preferred 2SLS baseline estimates, using frequencies of types (age, tenure, family/roommates, size) in the data over all years.

Next, in Figure 14 I plot rent coefficients (left panel) and bedroom fixed effects (right panel) for renting families of various sizes headed by 20-29 year olds. Because both family and non-family households are plotted in the left panel, we can see how the assumption of additive separability renders a constant gap $\beta_r^{f=1}$ between the two sets of coefficients. As expected, family households are more rent-sensitive than non-family households, and households of either type are more rent-sensitive when they are larger. Likewise, as expected, larger families have stronger preferences for the largest (3+ bedroom) units, in absolute terms and relative to mid-size (2 bedroom) units.

Given these results, rising housing costs (on average, or for the largest units) will disincentivize fertility. However, determining the magnitude of this effect requires using the full structural model. We turn to that next.

7.2 Decomposition

I measure the causal effect of rising rents on fertility by manipulating rents, solving the model, and calculating fertility. I provide results on the general fertility rate to facilitate comparisons with a reduced-form approach, the total number of children born, the decline in the fertility rate from the 2000s to 2010s, and the share of 20-29 year olds that have started a family. Because fertility is measured in children aged 0-9 in each decade, the x-axis labels on the following line plots refer

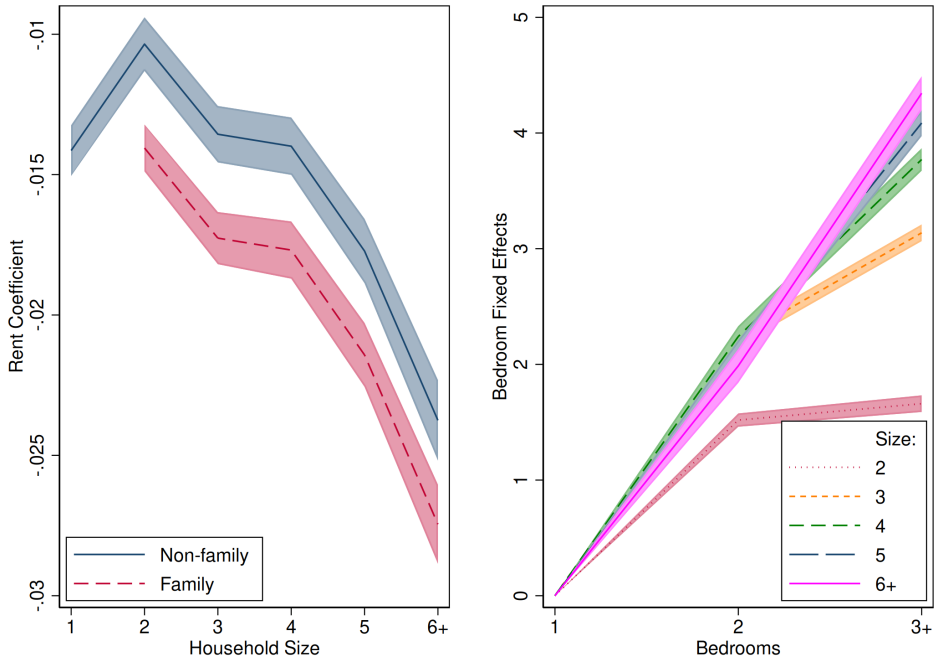


Figure 14: Parameters of interest for young renters. Large families are the most price-sensitive and place the highest value on living space.

to the decade ending in that year, and so the large decreases in the GFR, TFR, and children born between the last two periods are decreases from the 2000s to the 2010s[22]. I show that results are robust to an alternate normalization in Section A11.

In the top left panel of Figure 15, I plot the general fertility rate in the data, in the decomposition, and resulting from aggregating the linear predictions of the reduced-form specification in Section 3. The model and reduced-form agree in 2000, but the reduced form provides larger and larger yet predictions in 2010 and 2020, including a small increase in the GFR from the 2000s to the 2010s. Assuming the decomposition is correct, the bias in the reduced-form from sorting is substantial, equal to 31% and 68% of the true effect in the last two periods. If we took the reduced-form estimates at face value, it would imply that rising housing costs explain more than 100% of the decline in fertility from the 2000s to the 2010s. On the other hand, the decomposition provides a decrease in the GFR from 0.071 to 0.067 compared with a decrease in the data of 0.064 to 0.058. Because holding housing costs constant generates a decrease that is 58% as large, housing costs contribute 42% to the decline in the GFR from the 2000s to the 2010s.

Next is the total number of children born, in the top right panel of Figure 15. Using the reduced-

[22]This decrease is not as stark as in annual plots of fertility due to this aggregation, although it is based on the same underlying data since I re-weight the Census and ACS micro-data to match the decadal analogues of the age-specific birthrates published by the CDC.

form GFR results to calculate growing future cohorts of women (placing it in an apples-to-apples comparison with the decomposition, which also uses fertility in cohort size evolution) results in a larger bias in terms of births, exceeding 50 million in the final period where only 38 million are born in the data. The decomposition also predicts an increase in births from the 2000s to the 2010s, even as fertility rates decline, because of a growing cohort of reproductive-age women from higher past fertility. Across the final three periods, the decomposition predicts an increase of 13 million births, which is 11% of the total number of births in this period in the data.

In the bottom left panel of Figure 15 is the total fertility rate, which cannot be analyzed in the reduced-form due to data constraints. Holding housing costs fixed in the 2000s results in an increase in fertility substantially larger than that in the data, which likely occurred due to the strong economy at the time[23]. More interestingly, the decomposition features a decrease in the TFR of 0.11 compared with 0.21 in the data, implying that housing costs explain 51% of the decline in fertility from the 2000s to the 2010s.

Finally, the share of 20-29 year olds that are the head or the partner of the head of a household with children is in the bottom right panel. In both the data and the decomposition, this share declines in each period but especially in the last one. However, the total decline is only 10pp in the decomposition, compared with 17pp in the data. Housing costs are responsible for 7pp of the decline in family formation.

In subsequent decompositions depicted in Figure 16, I impose 1990 average rents (letting difference between sizes vary as in the data) and 1990 difference in rents (letting the average vary as in the data), in order to determine whether rising housing costs in general or the relative cost of large units caused declining fertility. The results show that rising average rents are the greater contributor to declining fertility, since reducing average rents down to 1990 levels does the most to increase both the total fertility rate and the number of children born, responsible for roughly 67-75%.

7.3 Counterfactuals

Finally, I turn to counterfactual simulations. As discussed in Section 6, I separately shift the supply curve the large 3+ bedroom units and small 1 bedroom units. Although I prefer to interpret these supply shifts as reductions in a distortionary regulatory tax, I define the magnitude of the shifts as if they come from subsidies of equal expenditure of 5% baseline rental expenditures (treating all households as renters), allowing the lower construction costs of small units to factor into the comparison. Then, the small unit policy may generate a large decrease in average rents, combining

[23]That the local peak of the US TFR occurred in 2007 is evidence that the expansion of housing finance in this period played a role in increasing fertility. Because tenure is a part of the living arrangement choice in the model which is not shifted in the decomposition, the decomposition is telling us that had housing costs (*i.e.*, mortgage rates) not increased at the same time that housing finance became easier to acquire, the fertility increase would have been even larger.

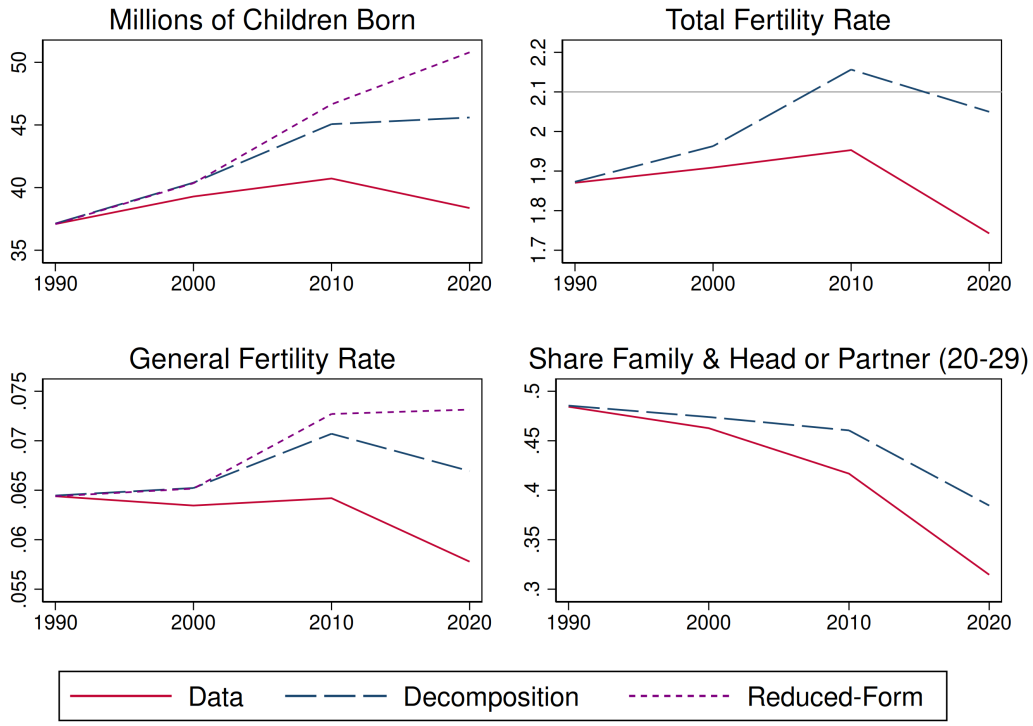


Figure 15: Results from holding rents at their 1990 levels. The blue long-dash line is the model decomposition. The purple short-dash line is from using earlier reduced-form estimates. Both include cohort size changes in calculations of children born.

with a greater dynamic effect from getting young adults out of their parents' houses (which shifts the long-run housing cost of fertility) to potentially surmount an increase in the relative rent of large units to generate a larger fertility increase than the large-unit policy.

However, as shown in Figure 17 I find that the large unit policy increases fertility by more, as both the number of births (top left) and total fertility rate (top right) is higher in every period in this scenario. In the bottom left, we verify that the small unit policy shifts average rents more, since the rent of 1 bedroom units is smaller by a lot compared with the large unit policy, the reduction in the rent of 2 bedroom units is about double, and the reduction in the rent of 3+ bedroom units is almost as large. However, this is insufficient to outweigh the impact of the large unit policy, which combines a smaller effect on average rents with a flattening of the gradient of rents with respect to size.

8 Conclusion

In this paper, I analyze the role of the housing market in declining fertility in the US. Motivated by the tendency for geographic mobility to generate selection into treatment in reduced-form analyses of the topic, I specify a structural model of joint housing and fertility choice. The model nests

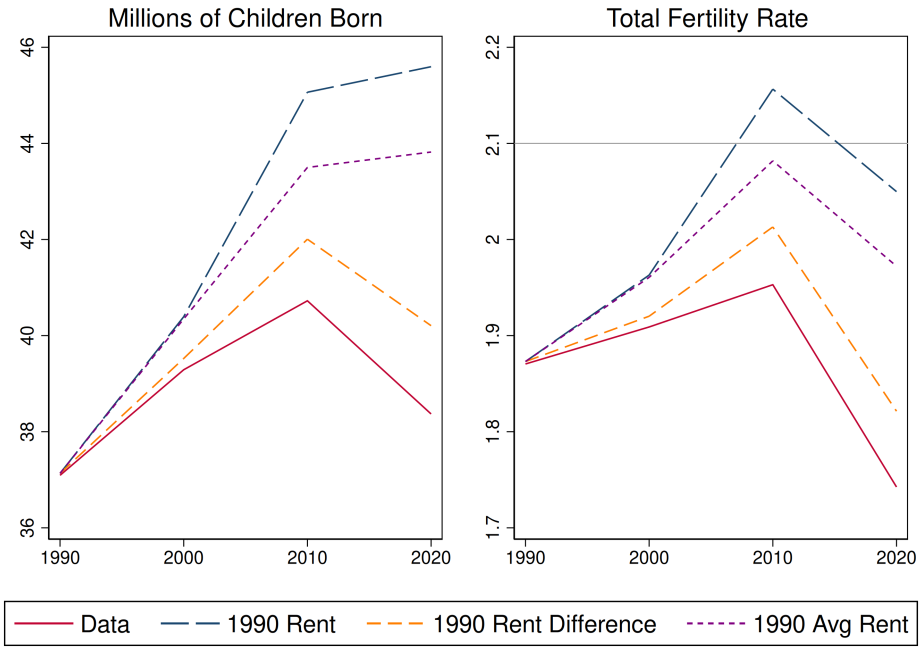


Figure 16: Results from imposing 1990 rents, 1990 rent differences, and 1990 average rents. The decomposition results (long-dash blue line) are driven by average rents (short-dash purple line), not rent differences (medium-dash orange line).

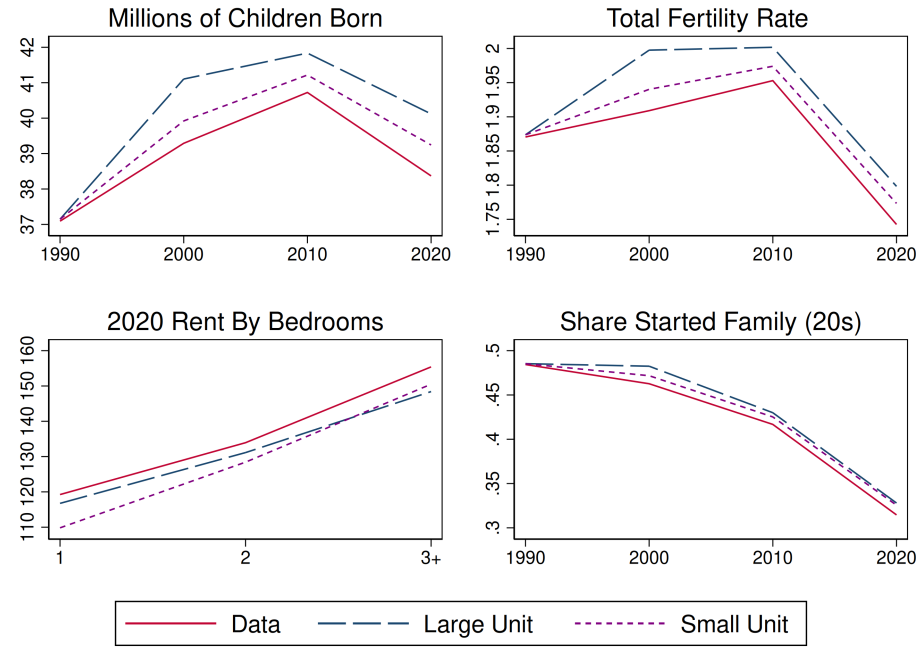


Figure 17: Counterfactual results of equal subsidies to the supply of large (3+ bedroom) and small (1 bedroom) units. The bottom left panel is not a time series panel; it plots 2020 rent over number of bedrooms.

a standard neighborhood sorting model – with an additional dimension of choice, the number of bedrooms – in a model of living arrangement choice, which chiefly includes fertility along with tenure and the option to live with parents, other family members, or unrelated roommates. The number of people in the household – be they the head, a partner, young children, adult children, or other family members – determines the preferences that the household will have over housing choice, generating a mechanism by which shifts in the characteristics of housing options translate into heterogeneous shifts in the value of having a large family. A quantity-quality model extended to include housing predicts, and housing demand estimates verify, that larger families will be more price-sensitive and have a stronger taste for living space, such that rising housing costs *or* a rising relative cost of large units disincentivize having children.

To estimate demand for neighborhoods that is heterogeneous over age, tenure, family/roommates, and household size despite the joint distribution of these variables being unobservable due to confidentiality, I utilize and extend IO "micro-moment" techniques. I show that not only is the necessary data for this housing demand model publicly available, it actually comprises *sufficient statistics* for model parameters under less restrictive assumptions than those imposed by the standard IO tools. In fact, the standard IO approach imposes an assumption that unobserved utility is homogeneous over demographics, which in some contexts will lead to misspecification. Conditional on the housing demand inclusive value, dynamic living arrangement (fertility) utility is chosen to fit non-geographic microdata exactly, so that all non-housing determinants of fertility – rising female education, changing culture, *etc.* – are held constant in decompositions and counterfactuals.

To find the causal effect of housing costs on fertility, I directly manipulate it in a partial-equilibrium counterfactual, and then measure fertility and living arrangement outcomes net of sorting and holding intrinsic living arrangement and fertility preferences constant. I find that rising housing costs since 1990 are responsible for 13 million (11%) children not being born, 51% of decrease in fertility from the 2000s to the 2010s, and a 7pp decrease in the share of 20-29 year olds that have started families. These effects are primarily driven by the general rise in housing costs rather than the relative cost of large units. I thus compare the fertility effects of shifting the supply of large and small units as if an equal aggregate subsidy had been applied to both, so that a larger number of lower-cost small units could potentially generate a larger impact on average rents and fertility (a reduction in distortionary local regulatory taxes is a more appealing interpretation and lower-cost policy, albeit one that does not provide two apples-to-apples supply shifts). However, I find that the large unit policy, which combines a smaller effect on average rents with a decrease in the rent of large units, has a 2.3 times larger impact on fertility. I conclude that although housing has potential as a lever in family policy, the focus must be on producing the housing that families want.

References

Ahlfeldt, Gabriel M. *et al.* (2025). "The Geography of Life: Evidence from Copenhagen". Working paper; abstract and slides available. URL: https://caterinasoto.github.io/PDF-website/Geography_of_life_presentation.pdf.

Albouy, David and R. Jason Faberman (2025). *Skills, Migration and Urban Amenities over the Life Cycle*. NBER Working Paper 33552. National Bureau of Economic Research. DOI: 10.3386/w33552.

Atalay, Kadir, Hsiu-Ling Li, and Shane Whelan (2021). "Housing Wealth, Fertility Intentions and Fertility". In: *Journal of Housing Economics* 53, p. 101737. DOI: 10.1016/j.jhe.2021.101737.

Athey, Susan and Guido W. Imbens (2007). "Discrete Choice Models with Multiple Unobserved Choice Characteristics". In: *International Economic Review* 48.4, pp. 1159–1192. DOI: 10.1111/j.1468-2354.2007.00458.x.

Autor, David, David Dorn, and Gordon Hanson (2019). "When Work Disappears: Manufacturing Decline and the Falling Marriage Market Value of Young Men". In: *American Economic Review: Insights* 1.2, pp. 161–178. DOI: 10.1257/aeri.20180010.

Baum-Snow, Nathaniel and Lu Han (2024). "The Microgeography of Housing Supply". In: *Journal of Political Economy* 132.6, pp. 1897–1946. DOI: 10.1086/728110.

Bayer, Patrick, Fernando Ferreira, and Robert McMillan (2007). "A Unified Framework for Measuring Preferences for Schools and Neighborhoods". In: *Journal of Political Economy* 115.4, pp. 588–638. DOI: 10.1086/522381.

Becker, Gary S. (1960). "An Economic Analysis of Fertility". In: *Demographic and Economic Change in Developed Countries*. Ed. by Universities-National Bureau Committee for Economic Research. NBER conference volume; institutional editor. Princeton, NJ: Princeton University Press, pp. 209–240. DOI: 10.7208/9780226042659-008.

Becker, Gary S. and H. Gregg Lewis (1973). "On the Interaction between the Quantity and Quality of Children". In: *Journal of Political Economy* 81.2, Part 2, S279–S288. DOI: 10.1086/260166.

Berry, Steven T. and Philip A. Haile (2010). *Identification in Differentiated Products Markets Using Market Level Data*. Cowles Foundation Discussion Paper 1744R. Cowles Foundation for Research in Economics. URL: https://cowles.yale.edu/publications/cfdp/cfdp-1744r.

— (2014). "Identification in Differentiated Products Markets Using Market Level Data". In: *Econometrica* 82.5, pp. 1749–1797. DOI: 10.3982/ECTA9027.

Berry, Steven T., James A. Levinsohn, and Ariel Pakes (July 1995). "Automobile Prices in Market Equilibrium". In: *Econometrica* 63.4, pp. 841–890. DOI: 10.2307/2171802.

Berry, Steven T., James A. Levinsohn, and Ariel Pakes (2004a). "Differentiated Products Demand Systems from a Combination of Micro and Market Share Data: The New Car Market". In: *Journal of Political Economy* 112.1, pp. 68–105. DOI: 10.1086/379939.

Berry, Steven T., Oliver B. Linton, and Ariel Pakes (2004b). "Limit Theorems for Estimating the Parameters of Differentiated Product Demand Systems". In: *The Review of Economic Studies* 71.3, pp. 613–654. DOI: 10.1111/j.1467-937X.2004.00298.x.

Black, Dan A. *et al.* (2013). "Are Children "Normal"?" In: *The Review of Economics and Statistics* 95.1, pp. 21–33. DOI: 10.1162/REST_a_00257.

Bloom, David E., Michael Kuhn, and Klaus Prettner (Aug. 2024). "Fertility in High-Income Countries: Trends, Patterns, Determinants, and Consequences". In: *Annual Review of Economics* 16.1, pp. 159–184. DOI: 10.1146/annurev-economics-081523-013750.

Borg, Mary O. (1989). "The Income–Fertility Relationship: Effect of the Net Price of a Child". In: *Demography* 26.2, pp. 301–310. DOI: 10.2307/2061527.

Börsch-Supan, Axel (1986). "Household Formation, Housing Prices, and Public Policy Impacts". In: *Journal of Public Economics* 30.2, pp. 145–164. DOI: 10.1016/0047-2727(86)90005-6.

Cesarini, David *et al.* (2023). *Fortunate Families? The Effects of Wealth on Marriage and Fertility*. NBER Working Paper 31039. National Bureau of Economic Research. DOI: 10.3386/w31039. URL: https://www.nber.org/papers/w31039.

Clark, Warren and Ana Ferrer (2019). "The Effect of Housing Prices on Fertility: Evidence from Canada". In: *Economics: The Open-Access, Open-Assessment E-Journal* 13.2019-38, pp. 1–31. DOI: 10.5018/economics-ejournal.ja.2019-38.

Clark, William A. V. (2012). "Do Women Delay Family Formation in Expensive Housing Markets?" In: *Demographic Research* 27, pp. 1–24. DOI: 10.4054/DemRes.2012.27.1.

Coeurdacier, Nicolas *et al.* (2022). "Fertility, Housing Costs and City Growth". Slides and abstract (no draft available). URL: https://floswald.github.io/project/fertility/.

Conlon, Christopher T. and Jeff Gortmaker (2020). "Best Practices for Differentiated Products Demand Estimation with PyBLP". In: *RAND Journal of Economics* 51.4, pp. 1108–1161. DOI: 10.1111/1756-2171.12352.

— (2023). *Incorporating Micro Data into Differentiated Products Demand Estimation with PyBLP*. NBER Working Paper 31605. Revised September 2024. National Bureau of Economic Research. DOI: 10.3386/w31605.

Cooper, Daniel and María José Luengo-Prado (2018). "Household Formation over Time: Evidence from Two Cohorts of Young Adults". In: *Journal of Housing Economics* 41, pp. 106–123. DOI: 10.1016/j.jhe.2018.05.004.

Couture, Victor and Jessie Handbury (2020). "Urban Revival in America". In: *Journal of Urban Economics* 119, p. 103267. DOI: 10.1016/j.jue.2020.103267.

Cumming, Fergus and Lisa J. Dettling (2024). "Monetary Policy and Birth Rates: The Effect of Mortgage Rate Pass-Through on Fertility". In: *The Review of Economic Studies* 91.1. First published online March 4, 2023, pp. 229–258. DOI: 10.1093/restud/rdad034.

Daysal, Nese Yıldız *et al.* (2021). "Home Prices, Fertility, and Early-Life Health Outcomes". In: *Journal of Public Economics* 198, p. 104395. DOI: 10.1016/j.jpubeco.2020.104395.

Dettling, Lisa J. and Melissa S. Kearney (2014). "House Prices and Birth Rates: The Impact of the Real Estate Market on the Decision to Have a Baby". In: *Journal of Public Economics* 110, pp. 82–100. DOI: 10.1016/j.jpubeco.2013.12.009.

— (2025). *Did the Modern Mortgage Set the Stage for the U.S. Baby Boom?* NBER Working Paper 33446. National Bureau of Economic Research. DOI: 10.3386/w33446.

Diamond, Rebecca (2016). "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980–2000". In: *American Economic Review* 106.3, pp. 479–524. DOI: 10.1257/aer.20131706.

Doepke, Matthias *et al.* (2023). "The Economics of Fertility: A New Era". In: *Handbook of the Economics of the Family*. Ed. by Shelly Lundberg and Alessandra Voena. Elsevier, pp. 1–86. DOI: 10.1016/bs.hefam.2023.01.003.

Driscoll, Anne K. and Brady E. Hamilton (Mar. 2025). *Effects of Age-specific Fertility Trends on Overall Fertility Trends: United States, 1990–2023*. Tech. rep. 3. Hyattsville, MD: National Center for Health Statistics. URL: https://www.cdc.gov/nchs/data/nvsr/nvsr74/nvsr74-3.pdf.

Ermisch, John (1999). "Prices, Parents, and Young People's Household Formation". In: *Journal of Urban Economics* 45.1, pp. 47–71. DOI: 10.1006/juec.1998.2081.

Ermisch, John and Pamela Di Salvo (1997). "The Economic Determinants of Young People's Household Formation". In: *Economica* 64.256, pp. 627–644. DOI: 10.1111/j.1468-0335.1997.tb00439.x.

Fazio, Dimas *et al.* (2025). *Housing and Fertility*. Research Paper 111/2025. Also circulated as CEPR DP19805 (2024) and Banco Central do Brasil Working Paper 612 (2024). FEB-RN. DOI: 10.2139/ssrn.5046571. URL: https://ssrn.com/abstract=5046571.

Figueiredo, Octávio, Paulo Guimarães, and Douglas Woodward (2003). "A Tractable Approach to the Firm Location Decision Problem". In: *The Review of Economics and Statistics* 85.1, pp. 201–204. DOI: 10.1162/003465303762687811.

Fleisher, Belton M. and Jr. Rhodes George F. (1979). "Fertility, Women's Wage Rates, and Labor Supply". In: *American Economic Review* 69.1, pp. 14–24. DOI: 10.2307/1802493.

Gandhi, Amit and Jean-François Houde (Oct. 2019). *Measuring Substitution Patterns in Differentiated-Products Industries*. NBER Working Paper 26375. Introduces the "Differentiation IVs"; revised

November 2020. National Bureau of Economic Research. DOI: 10.3386/w26375. URL: https://www.nber.org/papers/w26375.

Haurin, Donald R., Patric H. Hendershott, and Dongwook Kim (1993). "The Impact of Real Rents and Wages on Household Formation". In: *Journal of Urban Economics* 35.2, pp. 178–190. DOI: 10.1006/juec.1993.1009.

Heckman, James J. and James R. Walker (1990). "The Relationship Between Wages and Income and the Timing and Spacing of Births: Evidence from Swedish Longitudinal Data". In: *Econometrica* 58.6, pp. 1411–1441. DOI: 10.2307/2938322.

Kearney, Melissa S. and Riley Wilson (2018). "Male Earnings, Marriageable Men, and Nonmarital Fertility: Evidence from the Fracking Boom". In: *The Review of Economics and Statistics* 100.4, pp. 678–690. DOI: 10.1162/rest_a_00739.

Lafortune, Jeanne and Corinne Low (2023). "Collateralized Marriage? Homeownership and the Household Formation of Young Americans". In: *American Economic Journal: Applied Economics* 15.3, pp. 216–252. DOI: 10.1257/app.20200632.

Lesthaeghe, Ron (2014). "The Second Demographic Transition: A Review and Appraisal". In: *Proceedings of the National Academy of Sciences* 111.51, pp. 18112–18115. DOI: 10.1073/pnas.1420441111.

Liu, Ziqian and Yu Zhang (2024). "The Effect of House Prices on Fertility: Evidence from House Purchase Restrictions". Working paper, SSRN version. URL: https://ssrn.com/abstract=4808554.

Lovenheim, Michael F. and Kevin J. Mumford (2013). "Do Family Wealth Shocks Affect Fertility Choices? Evidence from the Housing Market". In: *The Review of Economics and Statistics* 95.2, pp. 464–475. DOI: 10.1162/REST_a_00331.

Martin, Joyce A. *et al.* (Dec. 2013). *Births: Final Data for 2012*. Tech. rep. 9. Hyattsville, MD: National Center for Health Statistics. URL: https://stacks.cdc.gov/view/cdc/26769.

Moreno-Maldonado, Ana and Clara Santamaría (2024). *Delayed Childbearing and Urban Revival: A Structural Approach*. CEPR Discussion Paper 19002. CEPR. URL: https://cepr.org/publications/dp19002.

Mullahy, John (1997). "Instrumental-Variable Estimation of Count Data Models: Applications to Models of Cigarette Smoking Behavior". In: *The Review of Economics and Statistics* 79.4, pp. 586–593. DOI: 10.1162/003465397557169.

National Center for Family & Marriage Research (2012). *Age Variation in the Divorce Rate, 1990 & 2010*. Family Profile FP-12-05. Bowling Green, OH: National Center for Family & Marriage Research, Bowling Green State University. URL: https://scholarworks.bgsu.edu/ncfmr_family_profiles/73/.

National Center for Family & Marriage Research (2014). *Age Variation in the Divorce Rate, 1990–2012*. Family Profile FP-14-16. Bowling Green, OH: National Center for Family & Marriage Research, Bowling Green State University. URL: https://scholarworks.bgsu.edu/ncfmr_family_profiles/151/.

— (2017). *Age Variation in the Divorce Rate, 1990 & 2015*. Family Profile FP-17-20. Bowling Green, OH: National Center for Family & Marriage Research, Bowling Green State University. URL: https://www.bgsu.edu/ncfmr/resources/data/family-profiles/wu-age-variation-divorce-rate-1990-2015-fp-17-20.html.

— (2019). *Age Variation in the Divorce Rate, 1990 & 2017*. Family Profile FP-19-13. Bowling Green, OH: National Center for Family & Marriage Research, Bowling Green State University. DOI: 10.25035/ncfmr/fp-19-13. URL: https://www.bgsu.edu/ncfmr/resources/data/family-profiles/allred-age-variation-divorce-rate-1990-2017-fp-19-13.html.

— (2021). *Age Variation in the Divorce Rate, 1990 & 2019*. Family Profile FP-21-16. Bowling Green, OH: National Center for Family & Marriage Research, Bowling Green State University. DOI: 10.25035/ncfmr/fp-21-16. URL: https://scholarworks.bgsu.edu/ncfmr_family_profiles/1229/.

— (2023). *Age Variation in the Divorce Rate, 1990 & 2021*. Family Profile FP-23-16. Bowling Green, OH: National Center for Family & Marriage Research, Bowling Green State University. DOI: 10.25035/ncfmr/fp-23-16. URL: https://www.bgsu.edu/ncfmr/resources/data/family-profiles/westrick-payne-lin-age-variation-divorce-rate-1990-2021-fp-23-16.html.

Paciorek, Andrew D. (2013). *The Long and the Short of Household Formation*. Finance and Economics Discussion Series 2013-26. Board of Governors of the Federal Reserve System. URL: https://www.federalreserve.gov/pubs/feds/2013/201326/201326pap.pdf.

Pan, Yinghao and Hao Yang (2022). "Impacts of Housing Booms on Fertility in China: A Perspective From Homeownership". In: *International Regional Science Review* 45.5, pp. 534–554. DOI: 10.1177/01600176211066472.

Petrin, Amil (2002). "Quantifying the Benefits of New Products: The Case of the Minivan". In: *Journal of Political Economy* 110.4, pp. 705–729. DOI: 10.1086/340779.

Ruggles, Steven *et al.* (2025). *IPUMS USA: Version 16.0*. dataset. Minneapolis, MN: IPUMS. DOI: 10.18128/D010.V16.0. URL: https://doi.org/10.18128/D010.V16.0.

Schroeder, Jonathan *et al.* (2025). *IPUMS National Historical Geographic Information System: Version 20.0*. dataset. Minneapolis, MN: IPUMS. DOI: 10.18128/D050.V20.0. URL: https://doi.org/10.18128/D050.V20.0.

Simon, Curtis and Robert Tamura (2009). "Do Higher Rents Discourage Fertility? Evidence from U.S. Cities, 1940–2000". In: *Regional Science and Urban Economics* 39.1, pp. 33–42. DOI: 10.1016/j.regsciurbeco.2008.06.003.

Tsai, Yung-Yu *et al.* (2022). *The Effect of Financial Resources on Fertility: Evidence from Administrative Data on Lottery Winners*. IEAS Working Paper 22-A007. Institute of Economics, Academia Sinica. URL: https://www.econ.sinica.edu.tw/~econ/pdfPaper/22-A007.pdf.

U.S. Census Bureau (2016). *2020 Census Research and Testing: Investigating the 2010 Undercount of Young Children—Examining the Coverage of Young Mothers*. Internal Research Memorandum. 2016.08. U.S. Census Bureau.

United Nations (2022). *World Population Prospects 2022: Summary of Results*. ST/ESA/SER.A/448. New York: United Nations, Department of Economic and Social Affairs, Population Division. URL: https://population.un.org/wpp/.

Wrenn, Douglas H., Junjian Yi, and Bo Zhang (2019). "House Prices and Marriage Entry in China". In: *Regional Science and Urban Economics* 74, pp. 118–130. DOI: 10.1016/j.regsciurbeco.2018.12.001.

Appendix Contents

A1 Example Data: Marginal Distributions and IPF **A-2**

A2 Radius Of Residential Search Regions **A-3**

A3 Sorting and SUTVA in the Reduced-Form **A-3**

A4 Quantity-Quality With Housing **A-5**
A4.1 Main Results A-5
A4.2 Weak Separability A-7

A5 Deriving Housing Demand **A-8**

A6 Micro-BLP Mis-specification **A-8**
A6.1 Theorem and Proof A-8
A6.2 Example A-10

A7 Sufficient Statistics: Choice Models **A-11**
A7.1 Setup A-12
A7.2 Theorem 1 . . . . . A-13
A7.3 Proof of Theorem 1 . . . . . A-13
A7.4 Corollary 1 . . . . . A-14
A7.5 Proof of Corollary 1 . . . . . A-14

A8 Sufficient Statistics: Other Models **A-14**
A8.1 Lemma 1 . . . . . A-14
A8.2 Proof of Lemma 1 . . . . . A-15
A8.3 Lemma 2 . . . . . A-16
A8.4 Proof of Lemma 2 . . . . . A-16
A8.5 Theorem 2 . . . . . A-17
A8.6 Proof of Theorem 2 . . . . . A-17
A8.7 Corollary 2 . . . . . A-18
A8.8 Proof of Corollary 2 . . . . . A-18
A8.9 Theorem 3 . . . . . A-18
A8.10 Proof of Theorem 3 . . . . . A-18

A9 Housing Demand Estimation Details **A-19**

A9.1 Agents To Households A-19

A9.2 Transforming Rent So Sufficiency Results Apply A-20

A9.3 Exposure Variable A-20

A10 Equilibrium Definition **A-21**

A11 Identification And Sensitivity **A-24**

A1 Example Data: Marginal Distributions and IPF

To better illustrate the data used, consider the following simple example. There are four census tracts – each of which belongs to one of two cities – and two age groups and two sexes. This is the unobserved joint distribution, depicted in Table 5's "Population" column. Because this joint distribution is unobserved, we would instead use three marginal distributions in Table 6: tract by age, tract by sex, and age by sex by city. These marginal distributions could be used directly in a bespoke maximum likelihood estimation routine, or they could first be combined into a rectangular estimate of the joint distribution using IPF (Table 5's right-most column) to be passed into a more standard maximum likelihood estimation routine.

Table 5: Example: Population by City, Census Tract, Age, Sex, and IPF Estimate

City	Census Tract	Age	Sex	Population	IPF Estimate
City A	Tract 101	Young	Male	520	514.432
City A	Tract 101	Young	Female	480	485.568
City A	Tract 101	Old	Male	310	315.568
City A	Tract 101	Old	Female	340	334.432
City A	Tract 102	Young	Male	450	455.568
City A	Tract 102	Young	Female	470	464.432
City A	Tract 102	Old	Male	290	284.432
City A	Tract 102	Old	Female	320	325.568
City B	Tract 201	Young	Male	610	607.979
City B	Tract 201	Young	Female	590	592.021
City B	Tract 201	Old	Male	330	332.021
City B	Tract 201	Old	Female	350	347.979
City B	Tract 202	Young	Male	500	502.021
City B	Tract 202	Young	Female	520	517.979
City B	Tract 202	Old	Male	310	307.979
City B	Tract 202	Old	Female	340	342.021

Table 6: Example: Marginal Population Distributions

(a) By Tract and Age

Tract	Age	Pop.
101	Young	1000
101	Old	650
102	Young	920
102	Old	610
201	Young	1200
201	Old	680
202	Young	1020
202	Old	650

(b) By Tract and Sex

Tract	Sex	Pop.
101	Male	830
101	Female	820
102	Male	740
102	Female	790
201	Male	940
201	Female	940
202	Male	810
202	Female	860

(c) By City, Age, and Sex

City	Age	Sex	Pop.
City A	Young	Male	970
City A	Young	Female	950
City A	Old	Male	600
City A	Old	Female	660
City B	Young	Male	1110
City B	Young	Female	1110
City B	Old	Male	640
City B	Old	Female	690

A2 Radius Of Residential Search Regions

Assume that with probability s movers search in a disc of radius R around their current home, and with probability $1-s$ search in a disc of radius R centered on some point in a different metro. The smallest inter-metro move is larger than the largest intra-metro move. Potential destinations are uniformly distributed in each disc.

Starting with the case with $s = 1$ (all within-metro moves), the probability of moving a distance less than or equal to d is simply the area of a disk with radius d divided by the total disk area:

$$Pr(D \leq d) = \pi d^2/\pi R^2 = (d/R)^2. \tag{A2.1}$$

Then we can use the median distance m to solve for the radius R:

$$Pr(D \leq m) = (m/R)^2 = 1/2 \implies R = m\sqrt{2}. \tag{A2.2}$$

When within-metro moves only occur with probability s, the ratio of the disk areas is now a conditional probability:

$$Pr(D \leq d) = Pr(D \leq d|s) \times s = \pi d^2/\pi R^2 s = (d/R)^2 s, \tag{A2.3}$$

and we can use median distance m and within-metro share s to solve for the radius R:

$$Pr(D \leq m) = (m/r)^2 s = 1/2 \implies R = m\sqrt{2s}. \tag{A2.4}$$

Given that the median move distance was stable between 16km and 24km from 1989 to 2019 (NAR, 2023), and that the intra-metro move share was around 0.75 and 0.8 over 2005-19 (ACS 2005, 2010, 2015, 2019), we arrive at a ballpark

$$R \in (16 \times \sqrt{2 \times 0.75},\ 24 \times \sqrt{2 \times 0.8}) \approx (19.6,\ 30.4). \tag{A2.5}$$

A3 Sorting and SUTVA in the Reduced-Form

To see the SUTVA violation, consider the following simplified model of joint sorting-fertility choice. Begin by positing that the utility of a joint location-fertility choice depends solely on that location's rents, a heterogeneous fertility constant, and an idiosyncratic EVT1 shock. Fertility is a binary decision $f \in \{0, 1\}$, and every individual i has a binary fertility type $\tau \in \{0, 1\}$, with $\alpha_1^1 > 0$ and fertility utility 0 otherwise. If the agent chooses fertility then they must also purchase an additional

unit of housing. The joint utility is:

$$u^i_{jf} = \alpha^{\tau(i)}_f + \beta(1+f)r_j + \epsilon^i_{jf}. \tag{A3.1}$$

The probability of choosing a (j, f) option is

$$P^\tau_{jf} = \frac{\exp(\alpha^\tau_f + \beta(1+f)r_j)}{\sum_k \exp(\alpha^\tau_0 + \beta r_k) + \exp(\alpha^\tau_1 + 2\beta r_k)}. \tag{A3.2}$$

The fertility rate for a given type and location is

$$F^\tau_j = \frac{P^\tau_{j1}}{P^\tau_{j0} + P^\tau_{j1}} = \frac{\exp(\alpha^\tau_1 + 2\beta r_j)}{\exp(\alpha^\tau_0 + \beta r_j) + \exp(\alpha^\tau_1 + 2\beta r_j)}, \tag{A3.3}$$

where notably the identical denominators cancel out of the logit choice probabilities. Although heterogeneous small-area fertility rates are generally unavailable, if they were observed and regressed on rents with a valid identification strategy (or, if aggregate fertility rates are used with a very strong assumption of homogeneity of fertility preferences), then at best the estimated coefficients will be a linear approximation to this non-linear function and large counterfactuals will be invalid.

Most likely, heterogeneous small-area fertility rates are unavailable, and the researcher would be forced to use aggregate small-area fertility rates. These are a weighted average of the heterogeneous fertility rates, where the weights are the proportions of either type that choose to settle in j:

$$F_j = W^1_j \frac{P^1_{j1}}{P^1_{j0} + P^1_{j1}} + (1 - W^1_j)\frac{P^0_{j1}}{P^0_{j0} + P^0_{j1}}. \tag{A3.4}$$

Although the denominators of the heterogeneous fertility rates cancel, they do not cancel out of the weights, since fertility types 0 and 1 have different choice probability denominators:

$$W^1_j = \frac{N^1_j}{N^0_j + N^1_j} = \frac{N\chi(P^1_{j0} + P^1_{j1})}{N(1-\chi)(P^0_{j0} + P^0_{j1}) + N\chi(P^1_{j0} + P^1_{j1})}, \tag{A3.5}$$

where N is the total number of agents and χ is the exogenous proportion of agents that are of fertility type 1. Because the denominators of P do not cancel out of W, W and thus F are functions of rents in all locations, and the reduced-form regression of F on r not only omits non-linearity but also violates SUTVA.

The sorting-induced SUTVA violation biases estimation. The partial derivative of F_j with respect to r_j picks up not just the causal impact of r_j on F^τ_j, but also a sorting effect as the composition of types changes through W. Being more fertile than the other and thus more rent-sensitive, the population of fertility type 1 choosing j decreases more as r_j increases, and j becomes

dominated by the less-fertile type 0. By implicitly making a structural assumption of no mobility, the reduced-form's treatment effects attribute causality to selection (and omit non-linearity) regardless of whatever identification strategy is used, with the magnitude of the bias in estimates and implied policy counterfactuals depending on the extent to which heterogeneous preferences for fertility are omitted.

A4 Quantity-Quality With Housing

A4.1 Main Results

The results that (1) larger families value living space more, and (2) larger families value the numeraire more (are more price-sensitive) emerge naturally from the following theoretical model, driven by the core idea that having more children pushes a family back onto a part of the utility curve where the marginal utility of living space per person and parental consumption are higher.

The key assumptions are (1) preferences are such that fertility is *weakly separable* from housing, child quality, and parental consumption, (2) goods are goods (have positive marginal utility), (3) goods are weakly normal (consumption non-decreasing with income), (4) goods have declining marginal utility, and (5) the elasticity of per-child housing and investment with respect to fertility is greater than -1. In addition, I impose standard assumptions (utility continuously differentiable, quasi-concave, *etc.*).

The consumer's problem is

$$\begin{aligned} &\max_{n,q,c,h} \mathcal{U}(n,q,c,h) \\ &\text{subject to } c + ni + nrh = y \\ &\text{and } q = f(i),\ f' > 0. \end{aligned} \tag{A4.1}$$

Parents choose the number of children n (real-valued for simplicity), their uniform quality q, parental consumption c, and housing per child h. Parental consumption, total investment in children, and total housing expenditure must satisfy a budget constraint[24]. Uniform investment in children i is converted into child quality q by a well-behaved function f. Notably, having a larger quantity of children increases the cost of increasing uniform child quality and housing per child – the exact mechanism highlighted by Becker and Lewis (1973) for child quality, extended to housing. We can also easily see how an increase in housing costs r affects fertility, as it increases the shadow price of fertility n, which is $i + rh$, reducing fertility via both the income and substitution effects.

[24]The core results would not change with a more explicit treatment of household size, for example by using a budget constraint $(c/2) + ni + (n+2)rh = y$.

Using weak separability[25], we can rewrite the problem as a nested problem[26], where n is taken as given in the inner problem and chosen in an outer problem on the basis of preferences for fertility *per se* and how it affects the value of the inner problem:

$$\begin{aligned} &\max_{n} U(n, v(n, r, y)) \\ &\text{where } v(n, r, y) = \max_{q,c,h} \; u(q, c, h) \\ &\text{subject to } c + ni + nrh = y \\ &\text{and } q = f(i), \; f' > 0. \end{aligned} \tag{A4.2}$$

This formulation permits a clean analysis of how a shift in fertility affects the other choice variables while still allowing fertility to be a choice variable itself.

Consider how an increase in n affects the inner problem. It raises the shadow price of h and q, and so the income and substitution effects combine to generate $h^*_n < 0$ and $q^*_n < 0$. Because h and q have declining marginal utility, an increase in n pushes h^* and q^* down onto a steeper part of the utility function, and so the marginal utilities of h and q increase.

Now consider c, where the income and substitution effects go in opposite directions. Write optimal consumption as income not spent on optimal child quality or optimal housing, then differentiate:

$$\begin{aligned} c^* &= y - ni^* - nrh^* \\ \implies c^*_n &= -(ni^*_n + i^*) - (nrh^*_n + rh^*) \\ &= -i^* \underbrace{(1 + \epsilon_{i,n})}_{+} - rh^* \underbrace{(1 + \epsilon_{h,n})}_{+} < 0. \end{aligned} \tag{A4.3}$$

Provided that both elasticities are greater than -1, parental consumption decreases with fertility, and likewise moves onto a steeper part of the utility curve.

The assumption that these elasticities are greater than -1 is simply an assumption that *total* expenditure on child quality and housing increase with the number of children.

$$\begin{aligned} d\ln(nrh) = d\ln n + d\ln r + d\ln h &= d\ln n\left(1 + \frac{d\ln h}{d\ln n}\right) > 0 \iff \epsilon_{h,n} > -1 \\ d\ln(ni) = d\ln n + d\ln i &= d\ln n\left(1 + \frac{d\ln i}{d\ln n}\right) > 0 \iff \epsilon_{i,n} > -1 \end{aligned} \tag{A4.4}$$

[25] Becker and Lewis's (1973) assumptions on preferences are only those that are necessary for an interior optimum to exist (instead focusing on the implication of quantity increasing the shadow price of quality and vice versa in the budget constraint), and so weak separability is not incompatible with their framework.

[26] I prove this in Section A4.2.

If this assumption were violated, then the family would respond to an additional child by decreasing the size of its house and its total expenditure on child quality. A family may reasonably keep these items constant in response to an additional child, but most would increase expenditure on these items such that on average these assumptions hold.

Because per-child investment and parental consumption have higher marginal utility at higher fertility, there is also a higher marginal utility of having resources available to spend on them – that is, a higher marginal utility of avoiding housing costs. Larger families are thus more sensitive to housing costs, and in the sorting model they will have a rent coefficient that is more negative. Additionally, they will consume a larger quantity of housing, spend more of their budget on housing, and have a larger marginal utility of housing.

A4.2 Weak Separability

In this section I show that weak separability allows the utility maximization problem to be written as a nested problem, where child quality, housing, and parental consumption are chosen in an inner problem conditional on the number of children, and fertility is chosen in an outside problem on the basis of preferences for fertility *per se* and how it affects the value of the inner problem.

In this context, the definition of weak separability is that for any bundles $z = (q, c, h),\ z' = (q', c', h')$

$$\mathcal{U}(n, z) \geq \mathcal{U}(n, z') \implies \mathcal{U}(n', z) \geq \mathcal{U}(n', z')\ \forall\ n,\ n'. \tag{A4.5}$$

That is, preferences for one bundle of (q, c, h) versus another do not depend on which n is chosen.

Then, we may define a utility function $u(q, c, h)$ (with cardinality following from some normalization) that does not depend on n, while n does affect the optimal choices of (q, c, h) through its influence on the budget constraint.

Now define a function

$$U(n, s) = \mathcal{U}(n, z_s) \tag{A4.6}$$

for any bundle z_s that generates $u(z_s) = s$. Because u is increasing in all arguments, U is increasing in s.

Finally, observe that for the inner problem's optimal bundle (q^*, c^*, h^*) and indirect utility $v(n, r, y) = u(q^*, c^*, h^*)$ we have

$$\mathcal{U}(n, q^*, c^*, h^*) = U(n, u(q^*, c^*, h^*)) = U(n, v(n, r, y)), \tag{A4.7}$$

and hence we may maximize $\mathcal{U}$ by choosing n to maximize U directly and through its impact on the inner problem's value.

A5 Deriving Housing Demand

Given the structure of the overall model, we can derive a standard residential choice model that uses the aggregated housing choice probabilities to estimate housing utility parameters.

To see this, write the housing choice probabilities (probability of choosing housing option h conditional on being type τ) as a function of the fundamental choice probabilities (probability of choosing housing-fertility option (g, h) conditional on being in state x):

$$P_h^\tau = \frac{Pr(h,\tau)}{Pr(\tau)} = \frac{\sum_{(x',g'):\tau(x',g')=\tau} P_{g'h}^{x'}}{\sum_{(x',g'):\tau(x',g')=\tau} \sum_{h'} P_{g'h'}^{x'}} \tag{A5.1}$$

The numerator of the fundamental choice probability in the numerator features a housing utility that by virtue of its index τ is common to all terms in the summation, and so can be factored out:

$$\sum_{(g',x'):\tau(g',x')=\tau} \exp(u_h^{\tau(g',x')} + \alpha_{g'}^{x'} + c_{g'}^{x'}) = \exp(u_h^\tau) \sum_{(g',x'):\tau(g',x')=\tau} \exp(\alpha_{g'}^{x'} + c_{g'}^{x'}). \tag{A5.2}$$

That utility is the sole term that is indexed by h, and so the choice probability is proportional to this exponentiated utility, exactly as in a standard logit model with choice over h alone:

$$P_h^\tau = \exp(u_h^\tau) K^\tau. \tag{A5.3}$$

Provided that a type fixed effect is included in the model, the parameters of u_h^τ can be estimated with maximum likelihood or by log-linearizing.

A6 Micro-BLP Mis-specification

A6.1 Theorem and Proof

A standard Micro-BLP model:

$$u_j^i = X_j' \beta^i + \xi_j + \epsilon_i^j, \quad \beta^i \sim N(\beta + \Pi y_i, \Sigma\Sigma'). \tag{A6.1}$$

The micro-BLP estimator uses market shares s, a number of micro-moments which often but not always grows with J, product characteristics X, and instruments Z. Notably, unobserved utility ξ is constrained to be constant over demographics.

Now consider an alternate model:

$$u_j^i = X_j' \beta^i + \xi_j^i + \epsilon_i^j, \quad \beta^i \sim N(\beta + \Pi y_i, \Sigma\Sigma'), \quad \xi_j^i = \xi_j + \Gamma_j y_i, \tag{A6.2}$$

that is, where unobservable utility is allowed to vary over demographics but follows some structure such that it is separately identified from ϵ_j^i. Or, more concisely, each fine demographic group g gets its own unobservable utility $\xi_j^{g(i)}$. The following result shows that the standard Micro-BLP model that constrains residuals to be constant over demographics generates testable restrictions on micro-moments when there are sufficiently many micro-moments, and that for generic data (that is, data that obeys accounting identities but which otherwise is not assumed to follow any particular pattern) the probability that these restrictions will be satisfied has measure zero.

Theorem: for J sufficiently large, fixed demographic specification, local identification, and a micro-moment set that grows with J, the Micro-BLP DGP generates observed data with probability zero and is mis-specified in the population.

Proof: Collect parameters and errors into $\theta = (\beta, \Pi, \Sigma, \xi)$. Collect market shares and micro-moments into a vector W. These parameters and errors generate the data in a system of equations

$$W = f(\theta). \tag{A6.3}$$

Let the dimensionality of θ be P, and the dimensionality of W (excluding linearly dependent equations) be D. The function f maps a point in the parameter space into the data space, and by virtue of the Micro-BLP model is continuously differentiable: $f : \Theta \subset \mathbb{R}^P \to \mathbb{R}^D$, $f \in C^1$. Given local identification, the Jacobian of f has full rank.

Given these assumptions, under the Implicit Function Theorem and/or the the Constant Rank Theorem, the image of f is a manifold of (at most) dimension P in $\mathbb{R}^D$. Parameters θ span (at most) $\mathbb{R}^P$, which the non-linear function f maps into a space that is locally euclidean of (at most) dimension P but globally non-euclidean: a manifold.

If $D > P$, the image of f has measure zero. The model generates the set of unrestricted data with probability zero, and so the model is mis-specified in the population.

When the number J of products j increases by one, a single ξ_j is appended to θ. However, since the number of micro-moments grows with J, a market share s_j and multiple micro-moments are appended to W. (Without mis-match between the market shares and micro-moments the market shares are sometimes functions of the micro-moments, but because the demographic shares that would be used to calculate the market shares from demographics are generally not passed into the estimation routine, they are non-redundant moments.) As J grows, D grows by more than P. Thus for a sufficiently large J^*, $D > P$, and Micro-BLP is mis-specified in the population. QED

The substantive assumptions in this proof are that (1) unobserved utility ξ is constant over demographics, and (2) the number of micro-moments increases with J. The first assumption is relaxed by the alternate DGP I specify. The second is true when the micro-moments are "partially dis-aggregated market shares" or market shares calculated over individual dimensions

of demographics $Pr(j|d^{\ell(i)})$, average consumer characteristics by product $E[y_i|j]$, or covariances between characteristics of the first and second choice $C(x_j, x_{k(-j)}|j,k)$. The second assumption is not true when the micro-moments are average product characteristics conditional on individual dimensions of demographics $E[x_j|d^{\ell(i)}]$ or covariances between product and consumer characteristics $C(x_j, y_i)$ are used, because these aggregate over products and do not grow with the number of products.

A6.2 Example

The threshold J^* need not be particularly large. Consider the following example with two dimensions of individual heterogeneity (a, s), no random coefficients, and the outside option included in the count J:

$$u_j^i = X_j'\beta^{(as)(i)} + \xi_j + \epsilon_j^i, \quad \beta_k^{as} = \beta_k + \beta_k^a + \beta_k^s \tag{A6.4}$$

The dimensionality of θ is $P = K(1 + A - 1 + S - 1) + J - 1$. For each characteristic, there is a baseline coefficient, and one age group and sex group each is normalized to have a zero coefficient. The outside option residual is normalized to zero.

With micro-moments of the form $P(j|d^{\ell(i)})$, the number of equations is $D = J - 1 + A(J - 1) + S(J - 1)$. There are aggregate and age- and sex-specific market shares, where the share for the outside option is a deterministic function of the others.

We then have

$$\begin{aligned} D > P &\iff J - 1 + A(J-1) + S(J-1) > K(A + S - 1) + J - 1 \\ &\iff AJ + SJ > AK + SK - K + A + S \\ &\iff J > \frac{(A+S)(K+1) - K}{A+S} \\ &\iff J > K + 1 - \frac{K}{A+S}. \end{aligned} \tag{A6.5}$$

For $K = 1$, accounting for integer constraints, we have $J^* = 2$: the model is mis-specified when there is a single inside product. Supposing there are two age groups and two sexes, there is a single inside market share, two age micro-moments, and two sex micro-moments, for a total of five moments. However, there is only the baseline coefficient, a coefficient for the non-default age group, and a coefficient for the non-default sex, plus one error, for a total of a four-element parameter-error vector.

Next, I depict this mis-specification numerically. I let $x_1 = 1$, $\beta = \beta^{a_1} = \beta^{s_1} = 1$, $(\xi_1^{a_0}, \xi_1^{a_1}, \xi_1^{s_0}, \xi_1^{s_1}) = (-1.5, -0.5, 0.5, 1.5)$, and each group equal to a quarter of the population. This generates an aggregate market share and four micro-moments that obey accounting identities.

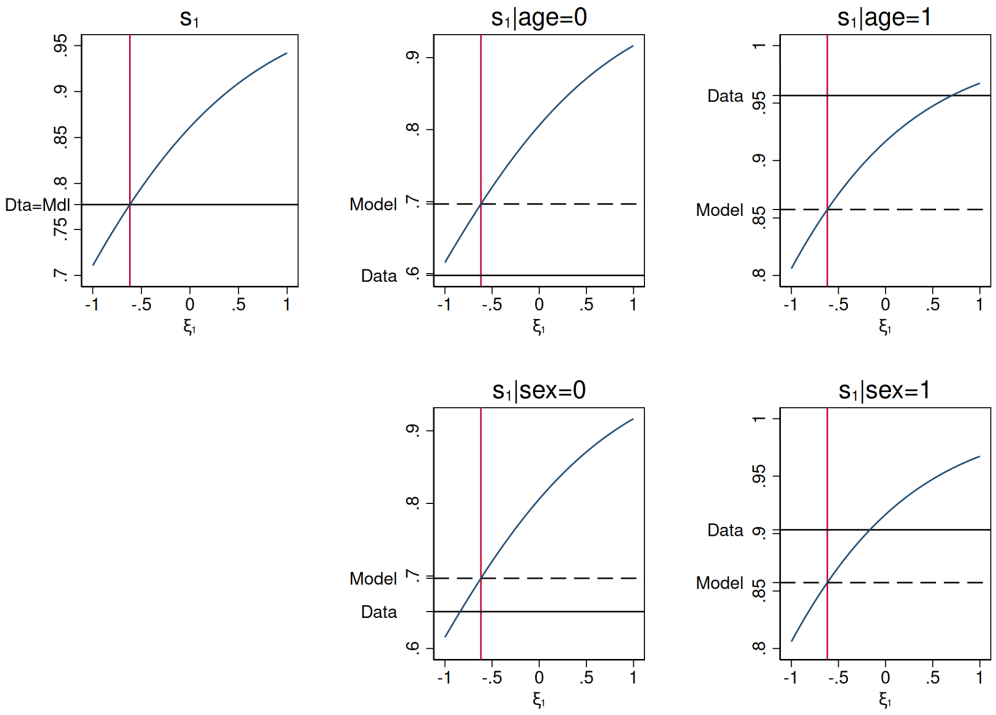


Figure 18

In the figure below, I let all coefficients be equal to their true values, and I let the single ξ_j fit the aggregate s_j. Because there are five moments plus a total of four parameters and errors, no single j is able to rationalize the data – the true parameters plus the single error can only rationalize a single moment at time.

More product characteristics in the model increases the number of parameters, increasing J^* at a rate less than 1. Adding dimensions of heterogeneity pushes the rate at which J^* increases with K up towards 1. Adding random coefficients would also increase the number of parameters, increasing J^*, but because the number of parameters does not grow with J there would still be finite J^*.

A7 Sufficient Statistics: Choice Models

Because the case of an arbitrary number of dimensions of heterogeneity each with an arbitrary number of dimensions requires complex and opaque notation, I illustrate these principles using the simplest possible setting.

A7.1 Setup

Assume a standard logit problem with a choice variable j (say, neighborhood of residence) and heterogeneity of preferences over two dimensions age a and sex s. There is a single characteristic x_j. Utility up to the EVT1 shock is

$$u_j^{as} = (\tilde{\beta} + \tilde{\beta}^a + \tilde{\beta}^s)x_j + \xi_j^{as} \implies u_j^{as} = (\beta^a + \beta^s)x_j + \xi_j^{as}, \tag{A7.1}$$

where in the first equation we write a baseline utility $\tilde{\beta}$ with one level each of $\tilde{\beta}^a$ and $\tilde{\beta}^s$ normalized to zero, and in the second the baseline utility is absorbed into one of the heterogeneity dimensions such that all coefficients in that dimension are shifted by $\tilde{\beta}$ and no level is normalized.

Then, the probability of an agent with demographics (a, s) choosing j is

$$P_j^{as} = \frac{\exp((\beta^a + \beta^s)x_j + \xi_j^{as})}{\sum_k \exp((\beta^a + \beta^s)x_k + \xi_k^{as})}, \tag{A7.2}$$

and the number of agents (the Poisson estimating equation) with demographics (a, s) choosing j is the product of the total number of (a, s) agents and that probability:

$$N_j^{as} = P_j^{as} N^{as} = \exp(\alpha^{as} + \beta^a x_j + \beta^s x_j + \xi_j^{as}) \tag{A7.3}$$

where

$$\exp(\alpha^{as}) = \frac{N^{as}}{\sum_j \exp(\beta^a x_j + \beta^s x_j + \xi_j^{as})}. \tag{A7.4}$$

As shown by Guimaraes et al (2003), after making this substitution the logit and Poisson likelihoods have the same maximizer $\hat{\beta}$, and so this logit model (with no endogeneity) can be estimated with PPML, with the demographic fixed effect α^{as} absorbing the choice probability denominator and the number of agents of group (a, s).

Let dropped indexes denote aggregation, so that the marginal distributions are

$$N_j^a = \sum_s N_j^{as}, \quad N_j^s = \sum_a N_j^{as}, \quad N^{as} = \sum_j N_j^{as}. \tag{A7.5}$$

These counts N_j^a, N_j^s, N^{as} are the counts of population (or households) over tract-age, tract-sex, and age-sex. The first two are published by the US Census Bureau as outputs from the Decennial Census and American Community Survey, whereas the last one is both published and derivable from microdata.

Remark: Note that (aside from the error term) the Poisson estimating equation contains additive terms varying at the $(a, s), (j, a), (j, s)$ levels, and no term (aside from the error term) varies at the

(j, a, s) level. This is the critical condition for the following results that assume that N^{as}, N^a_j, N^s_j are observed while N^{as}_j is not. Note that additive separability on preferences, $\beta^{as} = \beta^a + \beta^s$, generates this condition. If we let $u^{as}_j = \beta^{as} x_j + \xi^{as}_j$ then we would need to observe N^{as}_j.

We may also write N^{as}_j in terms of fixed effects (with a distinct error term since fixed effects explain at least as much as continuous explanatory variables):

$$N^{as}_j = \exp(\mu^s_j + \mu^a_j + \mu^{as} + \eta^{as}_j). \tag{A7.6}$$

Let $\hat{N}^{as}_j$ be the IPF-estimated count using the three two-dimensional marginal distributions:

$$\hat{N}^{as}_j = \text{IPF}(N^a_j, N^s_j, N^{as}). \tag{A7.7}$$

Let $\hat{\omega}^{as}_j$ be the error in the IPF-estimated count:

$$\hat{\omega}^{as}_j = N^{as}_j - \hat{N}^{as}_j. \tag{A7.8}$$

Because IPF matches the observed margins exactly, we have

$$\sum_j \hat{N}^{as}_j = N^{as}, \quad \sum_a \hat{N}^{as}_j = N^s_j, \quad \sum_s \hat{N}^{as}_j = N^a_j \tag{A7.9}$$

and so

$$\sum_j \hat{\omega}^{as}_j = 0 \; \forall (a, s), \quad \sum_a \hat{\omega}^{as}_j = 0 \; \forall \; (j, s), \quad \sum_s \hat{\omega}^{as}_j = 0 \; \forall \; (j, a). \tag{A7.10}$$

A7.2 Theorem 1

For any exponential conditional mean model with log-mean additively separable across the $(j, a), (j, s), (a, s)$ margins, $\hat{N}^{as}_j$ provides the same point estimates as N^{as}_j when used as a dependent variable.

A7.3 Proof of Theorem 1

Let $N^{as}_j = \exp(\lambda^{as} + \lambda^a_j + \lambda^s_j + \tau^{as}_j)$, where each λ may be FE or coefficient with explanatory variable.

For count outcome N^{as}_j the log-likelihood is

$$\mathcal{L} = \sum_j \sum_a \sum_s N_j^{as}(\hat{\lambda}^{as} + \hat{\lambda}_j^a + \hat{\lambda}_j^s) - \exp(\hat{\lambda}^{as} + \hat{\lambda}_j^a + \hat{\lambda}_j^s) \tag{A7.11}$$
$$= \sum_a \sum_s N^{as}\hat{\lambda}^{as} + \sum_j \sum_a N_j^a \hat{\lambda}_j^a + \sum_j \sum_s N_j^s \hat{\lambda}_j^s - \sum_j \sum_a \sum_s \exp(\hat{\lambda}^{as} + \hat{\lambda}_j^a + \hat{\lambda}_j^s),$$

where we have simply factored $\hat{\lambda}^{as}$ out of the summation over j so that we may use $\sum_j N_j^{as} = N^{as}$, and analogously for λ_j^a and λ_j^s, exploiting additive separability on preferences in the Poisson log-likelihood (one can use the exact same manipulation on the logit log-likelihood, which has the same maximizers). Thus the three margins N^{as}, N_j^a, N_j^s are sufficient statistics for $\hat{\lambda}$. QED

A7.4 Corollary 1

The two-step estimator – where (1) product fixed effects and heterogeneous slopes are estimated, and then (2) the product fixed effect is regressed on explanatory variables to recover the baseline utility coefficient – provides identical point estimates regardless of whether N_j^{as} or $\hat{N}_j^{as}$ is used as the dependent variable.

A7.5 Proof of Corollary 1

By the same argument as the proof of Theorem 1, N_j is sufficient for δ_j, and so the maximum likelihood estimator is the same regardless of whether N_j or microdata are used. Then, the second step 2SLS regression has the same δ_j, X, and Z regardless of whether N_j or microdata are used. Because N_j^a and N_j^s both imply N_j, the three two-way margins N_j^a, N_j^s, and N^{as} are sufficient for the Bayer *et al.* (2007) estimator.

A8 Sufficient Statistics: Other Models

Using the same setup as above.

A8.1 Lemma 1

The maximizer $\hat{\mu}$ for the Poisson log-likelihood

$$\mathcal{L} = \sum_a \sum_s \sum_j N_j^{as}(\hat{\mu}_j^a + \hat{\mu}_j^s + \hat{\mu}^{as}) - \exp(\hat{\mu}_j^a + \hat{\mu}_j^s + \hat{\mu}^{as}) \tag{A8.1}$$

satisfies

$$\exp(\hat{\mu}_j^a + \hat{\mu}_j^s + \hat{\mu}^{as}) = \hat{N}_j^{as} = \text{IPF}(N_j^a, N_j^s, N^{as}). \tag{A8.2}$$

A8.2 Proof of Lemma 1

The IPF algorithm consists of (1) initializing cells of the unobserved joint distribution with a positive count (*e.g.*, $1/N$), (2) cycling through marginal distributions, updating joint distribution counts proportionally so that at each stage a different marginal distribution fits exactly:

$$\begin{aligned}
\hat{N}_j^{as}(r) &= \hat{N}_j^{as}(r-1) \cdot \frac{N_j^a}{\sum_{s'} \hat{N}_j^{as'}(r-1)} \\
\hat{N}_j^{as}(r+1) &= \hat{N}_j^{as}(r) \quad \cdot \frac{N_j^s}{\sum_{a'} \hat{N}_j^{a's}(r)} \\
\hat{N}_j^{as}(r+2) &= \hat{N}_j^{as}(r+1) \cdot \frac{N^{as}}{\sum_{j'} \hat{N}_{j'}^{as}(r+1)},
\end{aligned} \tag{A8.3}$$

and (3) halting upon convergence below a specified tolerance.

Every IPF update strictly increases the log-likelihood, unless the margin already fit, in which case the increase is zero. Consider an update of the (j, a) margin, using the definition of the update and the fact that the grand totals always match the data:

$$\begin{aligned}
\mathcal{L}(r) - \mathcal{L}(r-1) &= \sum_a \sum_s \sum_j N_j^{as} (\ln \hat{N}_j^{as}(r) - \ln \hat{N}_j^{as}(r-1)) \\
&\quad + \sum_a \sum_s \sum_j \hat{N}_j^{as}(r) - \hat{N}_j^{as}(r-1) \\
&= \sum_a \sum_s \sum_j N_j^{as} \left(\ln \left(\frac{N_j^a}{\sum_{s'} \hat{N}_j^{as'}(r-1)} \cdot \hat{N}_j^{as}(r-1) \right) - \ln \hat{N}_j^{as}(r-1) \right) \\
&\quad + \hat{N}(r) - \hat{N}(r-1) \\
&= \sum_a \sum_s \sum_j N_j^{as} \left(\ln \left(\frac{N_j^a}{\sum_{s'} \hat{N}_j^{as'}(r-1)} \right) + \ln \hat{N}_j^{as}(r-1) - \ln \hat{N}_j^{as}(r-1) \right) \\
&\quad + N - N \\
&= \sum_a \sum_s \sum_j N_j^{as} \ln \left(\frac{N_j^a}{\sum_{s'} \hat{N}_j^{as'}(r-1)} \right)
\end{aligned} \tag{A8.4}$$

To see that this term is positive, consider the following function:

$$f(x) = x \ln x - x + 1, \quad f'(x) = \frac{x}{x} + \ln x - 1 = \ln x, \quad f''(x) = \frac{1}{x} > 0. \tag{A8.5}$$

This convex function achieves its unique minimum at $x = 1$, and therefore $x \ln x \geq 1 - x$ with equality only at $x = 1$. Now, let $x = N_j^a / \sum_{s'} \hat{N}_j^{as'}(r-1)$, rearrange the expression, and again use

the perfectly-fitting grand totals:

$$\begin{aligned} &\frac{N_j^a}{\sum_{s'} \hat{N}_j^{as'}(r-1)} \ln\left(\frac{N_j^a}{\sum_{s'} \hat{N}_j^{as'}(r-1)}\right) \geq 1 - \frac{N_j^a}{\sum_{s'} \hat{N}_j^{as'}(r-1)} \\ \implies & N_j^{as} \ln\left(\frac{N_j^a}{\sum_{s'} \hat{N}_j^{as'}(r-1)}\right) \geq \sum_{s'} \hat{N}_j^{as'}(r-1) - N_j^a = \hat{N}_j^a(r-1) - N_j^a \\ \implies & \sum_a \sum_j N_j^{as} \ln\left(\frac{N_j^a}{\sum_{s'} \hat{N}_j^{as'}(r-1)}\right) \geq \hat{N}(r-1) - N = 0 \,\forall\, s. \end{aligned} \tag{A8.6}$$

Every update strictly increases the log-likelihood, unless the margin already fits. When every margin fits, the algorithm has converged. At this point, the gradient of the log-likelihood is zero, and the likelihood is maximized.

A8.3 Lemma 2

$Z'\hat{\omega} = 0$, where Z is any stacked variable that varies (or sum of stacked variables that vary) at the $(j,a),(j,s),(a,s)$ level.

A8.4 Proof of Lemma 2

Collect all fixed effects into a single parameter vector θ. Define sets of vertical vectors A_j^a, B_j^s, C^{as} corresponding to the fixed effects $\mu_j^a, \mu_j^s, \mu^{as}$. They have as many rows as there are counts N_j^{as}, with elements equal to 1 if the indexes match the row index and 0 otherwise. Horizontally concatenate them to create D, which then has as many columns as fixed effects. Taking a single row of D, we then have

$$\hat{N}_j^{as} = \exp(D_j^{as} \cdot \hat{\theta}) \tag{A8.7}$$

and so the log-likelihood and gradient are

$$\begin{aligned} \mathcal{L} &= \sum_a \sum_s \sum_j N_j^{as} D_j^{as} \cdot \hat{\theta} - \exp(D_j^{as} \cdot \hat{\theta}) \\ \frac{\partial \mathcal{L}}{\partial \hat{\theta}} &= \sum_a \sum_s \sum_j N_j^{as} D_j^{as} - \hat{N}_j^{as} D_j^{as} \\ &= \sum_a \sum_s \sum_j \hat{\omega}_j^{as} D_j^{as}, \end{aligned} \tag{A8.8}$$

which is equal to zero at the optimum. This condition may be rewritten as

$$D'\hat{\omega} = \begin{pmatrix} A'\hat{\omega} \\ B'\hat{\omega} \\ C'\hat{\omega} \end{pmatrix} = \vec{0}. \tag{A8.9}$$

The errors in the counts are thus orthogonal to every column of D, a well-known property of the Poisson estimator. In this specific case, the errors are orthogonal to the indicators that define every fixed effect that varies at the $(j,a),(j,s),(a,s)$ level (every column of A, B, and C), and hence

$$Z'\hat{\omega} = \vec{0}, \tag{A8.10}$$

where Z is *any* variable that varies (or sum of variables that vary) at the $(j,a),(j,s),(a,s)$ level, due to collinearity of Z and D.

A8.5 Theorem 2

OLS and just-identified 2SLS models where the count is the dependent variable and the conditional mean is additively separable across the $(j,a),(j,s),(a,s)$ margins provide the same point estimates regardless of whether N_j^{as} or $\hat{N}_j^{as}$ is used.

A8.6 Proof of Theorem 2

Let each λ be either a fixed effect or a linear coefficient on an explanatory variable:

$$N_j^{as} = \lambda^{as} + \lambda_j^a + \lambda_j^s + \tau_j^{as} \implies \hat{N}_j^{as} = \lambda^{as} + \lambda_j^a + \lambda_j^s + \tau_j^{as} - \hat{\omega}_j^{as} \tag{A8.11}$$

Any endogenous explanatory variable either varies at one level with heterogeneous coefficients, or varies at two levels with a single coefficient. Thus, the instrument must also vary at a maximum of two levels, and the first stage conditional mean is also additively separable across the $(j,a),(j,s),(a,s)$ margins.

The difference between the parameter estimates is (letting $Z = X$ in the case of OLS)

$$\hat{\beta}^N - \hat{\beta}^{\hat{N}} = (Z'X)^{-1}Z'\hat{\omega}. \tag{A8.12}$$

By Lemma 2, $Z'\hat{\omega} = 0$ because in the cases of OLS and 2SLS Z is additively separable across the $(j,a),(j,s),(a,s)$ margins, and so $\hat{\beta}^N = \hat{\beta}^{\hat{N}}$.

A8.7 Corollary 2

models where the log count is the dependent variable and the conditional mean is additively separable across the $(j, a), (j, s), (a, s)$ margins provide the same point estimates regardless of whether N_j^{as} or $\hat{N}_j^{as}$ is used, provided that the model is estimated with PPML.

A8.8 Proof of Corollary 2

Proof: Follows directly from Theorem 1.

Remark: Often in reduced-form contexts, models with log counts as dependent variable are of greater interest than the untransformed counts. This corollary states that $\hat{N}_j^{as}$ can be used in these cases, but with PPML rather than OLS, since $\ln \hat{N}_j^{as} = \ln(N_j^{as} - \hat{\omega}_j^{as})$ is not additive in $\hat{\omega}_j^{as}$. Such a model is econometrically identical to that discussed in Theorem 1, dropping only the choice model interpretation.

$$\ln(N_j^{as}) = \lambda^{as} + \lambda_j^a + \lambda_j^s + \tau_j^{as} \implies \hat{N}_j^{as} = \exp(\lambda^{as} + \lambda_j^a + \lambda_j^s + \tau_j^{as}) - \hat{\omega}_j^{as} \tag{A8.13}$$

A8.9 Theorem 3

OLS and just-identified 2SLS models where the dependent variable is a count as a share of an observed margin, and where the conditional mean is in the space of the observed margin, provide the same point estimates regardless of whether N_j^{as} or $\hat{N}_j^{as}$ is used.

A8.10 Proof of Theorem 3

Without loss, divide the counts by N_j^a. We cannot include in the conditional mean any terms indexed by s. λ_j^a may be a fixed effect, a heterogeneous coefficient on an explanatory variable that varies over one dimension, or a coefficient on an explanatory variable that varies over two dimensions. Thus, in the case of 2SLS the instrument varies at 2 levels at most.

$$\frac{N_j^{as}}{N_j^a} = \lambda_j^a + \tau_j^{as} \implies \frac{\hat{N}_j^{as}}{N_j^a} = \lambda_j^a + \tau_j^{as} - \frac{\hat{\omega}_j^{as}}{N_j^a} \tag{A8.14}$$

The difference between the parameter estimates is (letting $Z = X$ in the case of OLS)

$$\hat{\beta}^N - \hat{\beta}^{\hat{N}} = (Z'X)^{-1} Z' \left(\frac{\hat{\omega}}{N_J^A} \right). \tag{A8.15}$$

In summation notation,

$$Z'\left(\frac{\hat{\omega}}{N_J^A}\right) = \sum_j \sum_a \sum_s z_j^a \frac{\hat{\omega}_j^{as}}{N_j^a} = \sum_j \sum_a \frac{z_j^a}{N_j^a} \underbrace{\sum_s \hat{\omega}_j^{as}}_{0} = 0 \tag{A8.16}$$

and so $\hat{\beta}^N = \hat{\beta}^{\hat{N}}$. This argument holds when z_j^a is replaced by an instrument that varies at a lower level, a continuous explanatory variable in the structural equation, or an indicator that defines fixed effects in the structural equation.

Remark: We restrict the conditional mean to (j, a) space because the argument fails once a term indexed by s is included. To see this, consider the following model:

$$\frac{\hat{N}_j^{as}}{N_j^a} = \beta^s x_j + \tau_j^{as} - \frac{\hat{\omega}_j^{as}}{N_j^a}. \tag{A8.17}$$

But now, we cannot isolate $\hat{\omega}_j^{as}$ in a summation over s, or j or a, because each index appears on at least one other term:

$$\sum_j \sum_a \sum_s \beta^s x_j \frac{\hat{\omega}_j^{as}}{N_j^a} \neq 0. \tag{A8.18}$$

A9 Housing Demand Estimation Details

A9.1 Agents To Households

Fundamentally, we wish to understand the behavior of the adults who comprise the households, and so revealed-preference estimators require data on counts of adults. With the data used in housing demand estimation, I assume that family households are comprised of two household decisionmakers, with the other members being children or extras:

$$N_{jbt}^{aefs} = \begin{cases} \hat{H}_{jbt}^{ae0s} \times s & \text{if} \quad f = 0 \\ \hat{H}_{jbt}^{ae1s} \times 2 & \text{if} \quad f = 1. \end{cases} \tag{A9.1}$$

This is consistent with letting housing utility be zero for adult children and extras: they unilaterally add themselves to the host household (increasing its size and thus housing preferences) but do not participate in the housing decision – and since they are distinguished only in the microdata we do not specifically observe the housing decisions of their hosts anyway. I also assume that every person in a non-family household is a decision-making adult. Because not every input marginal distribution is indexed by f and s, this transformation can only be applied after IPF estimation.

A9.2 Transforming Rent So Sufficiency Results Apply

Substituting Equation 4.11 into Equation 4.10 and inspecting the rent term, we obtain

$$\beta_r^\tau r_{jbt} = \beta r_{jbt} + \beta^a r_{jbt} + \beta^e r_{jbt} + \beta^f r_{jbt} + \beta^s r_{jbt}. \tag{A9.2}$$

Applying the heuristic for sufficiency, we must observe the counts $N_{jbt}, N^a_{jbt}, N^e_{jbt}, N^f_{jbt}, N^s_{jbt}$. Unfortunately, we do not observe $N^a_{jbt}, N^f_{jbt}, N^s_{jbt}$, because the only observed tract-level marginal distribution over j, b, and some other dimension is N^e_{jbt}.

I address this problem by constructing a rent variable that has a tract component and a PUMA-bedrooms component using regression:

$$r_{jbt} = \gamma_{jt} + \gamma_{p_t(j),b,t} + \nu_{jbt} \implies \hat{r}_{j,b,t} = \hat{\gamma}_{jt} + \hat{\gamma}_{p_t(j),b,t}. \tag{A9.3}$$

Using this term in housing utility:

$$\begin{aligned}\beta_r^\tau \hat{r}_{jbt} = {} & \beta_r \hat{\gamma}_{jt} + \beta_r \hat{\gamma}_{p_t(j),b,t} + \beta_r^a \hat{\gamma}_{jt} + \beta_r^a \hat{\gamma}_{p_t(j),b,t} \\ & + \beta_r^e \hat{\gamma}_{jt} + \beta_r^e \hat{\gamma}_{p_t(j),b,t} + \beta_r^f \hat{\gamma}_{jt} + \beta_r^f \hat{\gamma}_{p_t(j),b,t} + \beta_r^s \hat{\gamma}_{jt} + \beta_r^s \hat{\gamma}_{p_t(j),b,t},\end{aligned} \tag{A9.4}$$

and so the marginal distributions of tract with demographics, and of PUMA with everything (except tract) are sufficient. This is why I use PUMA as the higher-level geography in IPF estimation – although only city (larger than PUMA) is required for the city-type-year fixed effects, PUMA fixed effects fit the rent regression above better and allow more geographic-bedrooms variation in rents to be retained. Hereafter, I refer to $\hat{r}_{jbt}$ with r_{jbt} to keep notation simple.

A9.3 Exposure Variable

A location may have a large number of agents (or a high market share) because it provides a high utility, or simply because it is large. Therefore, some size adjustment is necessary for location choice problems. Given two locations of equivalent utility but different sizes, we would expect both to have the same density and for the larger one to have proportionally more people, motivating treating counts as proportional to size – a Poisson exposure variable. Area-based measures are difficult to work with since one must take a stand on what sorts of land qualify – much land in the world is inhospitable, which could be addressed with more amenities data or a measure of "habitable" or "buildable" land, but this poses a large burden. Housing-unit-based measures can pose their own problems, since they are a measure of quantity that is determined in a long run equilibrium of supply and demand, but they can be satisfying under the assumption that supply is fixed in the short run.

A10 Equilibrium Definition

A partial equilibrium is an input vector of rent shifts and a vector of value functions that satisfies the following conditions.

Housing utility is a linear function of rents.

$$u^{\tau}_{jbt} = \beta^{\tau}_{r} r_{jbt} + \beta^{\tau}_{b} + \beta^{\tau}_{m} m_{jt} + \beta^{\tau}_{x} x_{j} + \phi^{\tau}_{c(j),t} + \xi^{\tau}_{jbt}. \tag{4.10}$$

Value functions are the sum of housing utility, living arrangement utility, and the continuation value, where the continuation value is the discounted expectation of future values, which depends on the probability of transition to states given the current state and living arrangement choice and the inclusive value of each state.

$$v^{x}_{gh} = \underbrace{u^{\tau(g,x)}_{h} + \alpha^{x}_{g}}_{\text{flow utility}} + \underbrace{\delta \mathbb{E}\left[\underbrace{\sum_{x'} \pi(x' \mid g, x)}_{\text{state transition}} \times \underbrace{\ln\left(\sum_{h'} \sum_{g' \in \mathcal{G}(x')} \exp(v^{x'}_{g'h'}) \right)}_{\text{inclusive value}} \,\middle|\, \underbrace{W_{t(x)}}_{\text{info}_t} \right]}_{\text{continuation value} \equiv c^{x}_{g}}. \tag{4.3}$$

Living arrangement choice probabilities are the sum of joint choice probabilities, and a function of living arrangement utilities, living arrangement continuation values, and housing inclusive values.

$$P^{x}_{g} = \sum_{h} P^{h}_{gx} = \frac{\exp(\alpha^{x}_{g} + c^{x}_{g}) \sum_{h} \exp(u^{\tau(g,x)}_{h})}{\sum_{h'} \sum_{g' \in \mathcal{G}(x)} \exp(\alpha^{x}_{g'} + c_{g'} + u^{\tau(g',x)}_{h'})}. \tag{5.5}$$

The distribution of agents over states in the next period is given by an accounting relationship that depends on the distribution of agents over states in the current period, current period living arrangement choice probabilities, an endogenous increase from fertility, and exogenous changes from mortality and net migration.

$$N^{\tilde{x},t+1} = \sum_{\tilde{x}'} \sum_{g \in \mathcal{G}(x)} \pi(\tilde{x}, t+1 | \tilde{z}', x', g, t) N^{\tilde{x}'t} P^{\tilde{x}'t}_{g} + \Delta^{\tilde{x},t+1}_{\text{fertility}} + \Delta^{\tilde{x},t+1}_{\text{mortality}} + \Delta^{\tilde{x},t+1}_{\text{net migration}}. \tag{5.11}$$

Agents believe that the next period distribution of agents over states is deterministically the point estimate from an autoregression with fixed effects and time trends.

$$N^{\tilde{x},t+1} \approx \tilde{N}^{\tilde{x},t+1} = \hat{\psi}_{0} N^{\tilde{x}t} + \hat{\psi}^{\tilde{x}}_{1} + \hat{\psi}^{\tilde{x}}_{2} t, \tag{5.13}$$

The distribution of children aged 10-19 (and 20-29 analogously) in a household in the next period is given by a binomial distribution where the number of trials is the number of children and the probability is the endogenous probability of children of the relevant state choosing to live with their parents.

$$Pr(n_{10-19} = k|x, n_{0-9,t-1}) = \left(\frac{n_{0-9,t-1}!}{k!(n_{0-9,t-1}-k)!}\right)\left(P_{\text{par}}^{10-19,t,\text{imm}(x)}\right)^{k} \times \left(1 - P_{\text{par}}^{10-19,t,\text{imm}(x)}\right)^{(n_{0-9,t-1})-k}, \tag{4.5}$$

The total number of extras is the sum of all agents who choose to be extras.

$$N_{\text{extras}}^{t} = \sum_{\tilde{x}} N^{\tilde{x}t} P_{\text{extra}}^{\tilde{x}t}. \tag{4.6}$$

Agent states (excluding year) are mapped to an "extras group" w, which receives a constant share of all extras that comes directly from the data.

$$s^{w(\tilde{x})} = \frac{\sum_t \sum_{\tilde{x}':w(\tilde{x}')=w(\tilde{x})} N^{\tilde{x}'t} e(x')}{\sum_t \sum_{\tilde{x}''} N^{\tilde{x}''t} e(x'')}. \tag{4.7}$$

The number of extras that an "extras group" receives is its allocated share of the total number of extras.

$$N_{\text{extras}}^{w(\tilde{x}),t} = s^{w(\tilde{x})} N_{\text{extras}}^{t}. \tag{4.8}$$

The distribution of extras that an agent hosts is a binomial distribution where the number of trials is the number of extras that an "extras group" receives and the probability is the reciprocal of the number of agents in the "extras group" (every agent in the group has an equal probability).

$$Pr(\text{extras} = e|x) = \left(\frac{N_{\text{extras}}^{w(\tilde{x}),t}!}{e!\left(N_{\text{extras}}^{w(\tilde{x}),t} - e\right)!}\right)\left(\frac{1}{N^{w(\tilde{x}),t}}\right)^{e}\left(1 - \frac{1}{N^{w(\tilde{x}),t}}\right)^{N_{\text{extras}}^{w(\tilde{x}),t}-e}. \tag{4.9}$$

Agents form expectations of future values using an autoregression with fixed effects and time trends.

$$v_g^{\tilde{x},t+1} = \rho_0 v_g^{\tilde{x},t} + \rho_{g,1}^{\tilde{x}} + \rho_{g,2}^{\tilde{x}} t + e_g^{\tilde{x},t+1} \tag{5.8}$$

Conditional on a vector of value function error terms d, the distribution of children aged 10-19 (and 20-29 analogously) in a household in the next period is given by a binomial distribution where the number of trials is the number of children and the probability is the endogenous probability of

children of the relevant state choosing to live with their parents which depends on the draw.

$$Pr(n_{10-19} = k|x, n_{0-9,t-1}, d) = \left(\frac{n_{0-9,t-1}!}{k!(n_{0-9,t-1}-k)!}\right)\left(P_{\text{par}}^{10-19,t,\text{imm}(x),d}\right)^{k} \times \left(1 - P_{\text{par}}^{10-19,t,\text{imm}(x),d}\right)^{(n_{0-9,t-1})-k}. \tag{5.10}$$

Conditional on a vector of value function error terms d, the number of extras that an extras group believes they will receive depends on their prediction of the distribution of agents over states and the endogenous probabilities of agents choosing to be extras which depends on the draw.

$$\tilde{N}_{\text{extras}}^{\tilde{x},t+1,d} = s^{w(\tilde{x})} \sum_{\tilde{x}'} \tilde{N}^{\tilde{x}',t+1} P_{\text{extra}}^{\tilde{x},t+1,d}. \tag{5.14}$$

Conditional on a vector of value function error terms d, agents believe that the distribution of extras that they will host is a binomial distribution where the number of trials is the number of extras they believe their group will receive (which depends on the draw) and the probability parameter is the reciprocal of the number of agents they believe will be in the extras group.

$$Pr(\text{extras} = e|\tilde{x}, t+1, d) = \left(\frac{\tilde{N}_{\text{extras}}^{w(\tilde{x}),t+1,d}!}{e!(\tilde{N}_{\text{extras}}^{w(\tilde{x}),t+1,d}-e)!}\right)\left(\frac{1}{\tilde{N}^{w(\tilde{x}),t+1}}\right)^{e}\left(1 - \frac{1}{\tilde{N}^{w(\tilde{x}),t+1}}\right)^{\tilde{N}_{\text{extras}}^{w(\tilde{x}),t+1,d}-e}. \tag{5.15}$$

An equilibrium is an input vector of supply shifts and a vector of value functions that satisfy the conditions above, and the following conditions.

Housing choice probabilities are a logit function of housing utilities.

$$P_h^{\tau} = \exp(u_h^{\tau})K^{\tau}. \tag{A5.3}$$

Right to left (recalling $h = (j, b)$). The number of agents in a housing type is the sum over agents who make living arrangement choices in states which in combination render them that housing type. The number of agents who are of a given housing type and choose a given housing option is the product of the number of agents of that housing type and the probability that agents of that housing type choosing that housing option. Roommate households are completely comprised of decisionmaking agents while family households are assumed to have two decisionmaking agents (even as they may host others who do not participate in the housing decision) such that the number of households can be obtained from inverting those functions. The total number of households of a given housing type choosing a housing option is a sum of the number of households of a given housing type who choose that option.

$$H_{jbt} = \sum_{\tau} H_{jbt}^{\tau}, \quad H_{jbt}^{\tau} = \begin{cases} N_{jbt}^{\tau}/s(\tau) & \text{if} \quad f(\tau) = 0 \\ N_{jbt}^{\tau}/2 & \text{if} \quad f(\tau) = 1, \end{cases} \quad N_{jbt}^{\tau} = P_{jbt}^{\tau} N^{\tau}, \quad N^{\tau} = \sum_{(g,x):\tau(g,x)=\tau} P_g^x N^x \tag{6.2}$$

Housing supply has a constant elasticity and heterogeneous intercepts.

$$\ln r_{jbt} = \gamma_{jb} + \gamma_{jt} + \gamma_{bt} + \gamma_j^{-1} \ln H_{jbt} + \epsilon_{jbt}. \tag{6.1}$$

A11 Identification And Sensitivity

The preferred normalization is that the normalized choice is living alone when the state is non-family, having zero children and not buying when the state is renter family, and having zero children when the state is owner family. The alternate normalization used for sensitivity analysis is that the normalized choice is living in a two-person roommate households when the state is non-family, having one child and not buying when the state is renter family, and having one child when the state is owner family. Family households that are unable to have children due to having lost a partner or being too old continue to have a normalized choice of zero children. The preferred normalization is preferred because it most closely matches intuitive notions of a default choice.

The alternative normalization provides qualitatively similar but numerically smaller impacts of housing on fertility. It continues to support the conclusion that a large-unit policy is better for fertility. Because it provides very similar results for the share of young people that have started a family and for rents, it seems likely that the differences arise from the shift in default from zero to one child, rather than from living alone to with a single roommate for agents that have not started a family or some inherent instability. This is further evidence that the preferred normalization is more appropriate: zero children is much more appealing as an "outside option" than having one child.

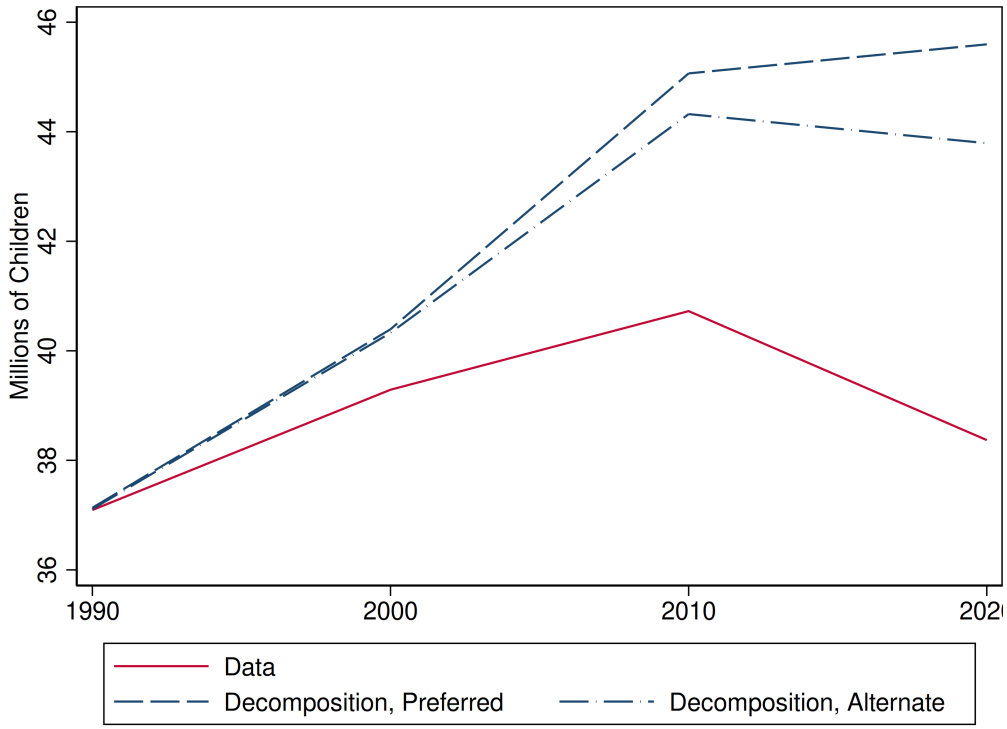


Figure 19: Decomposition: Children

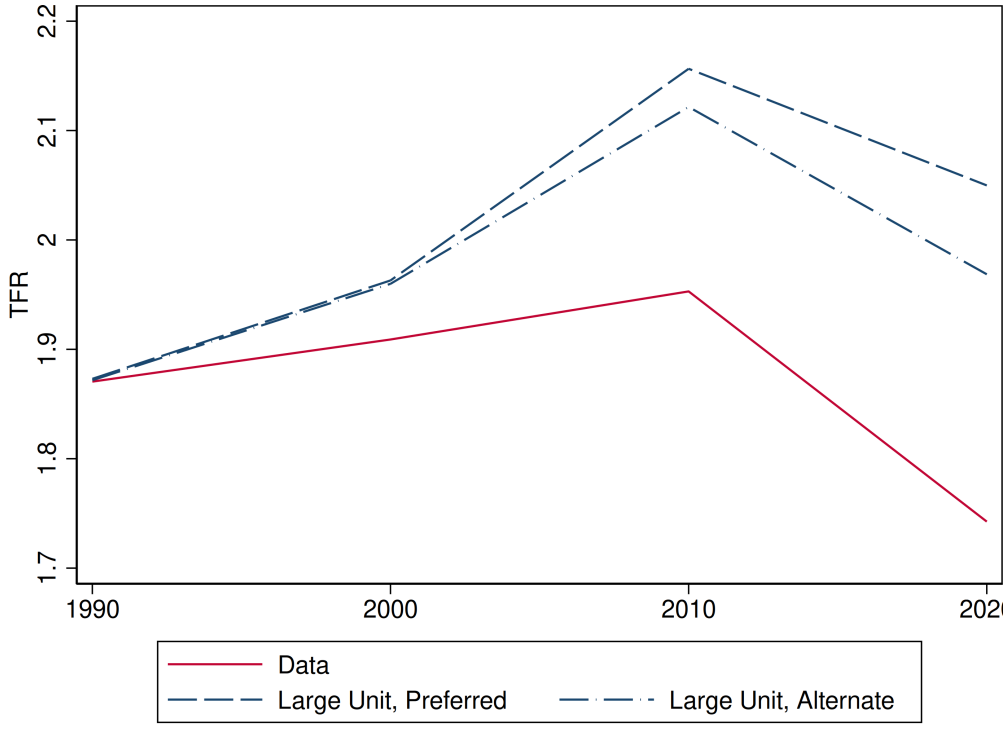


Figure 20: Decomposition: TFR

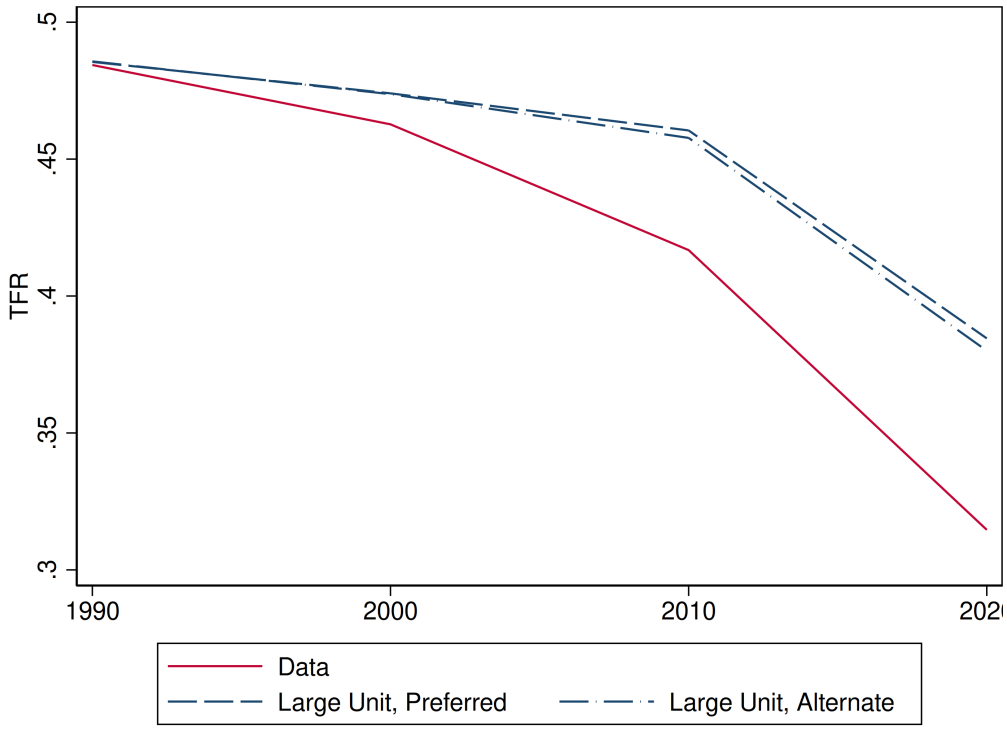


Figure 21: Decomposition: Share Family, 20-29

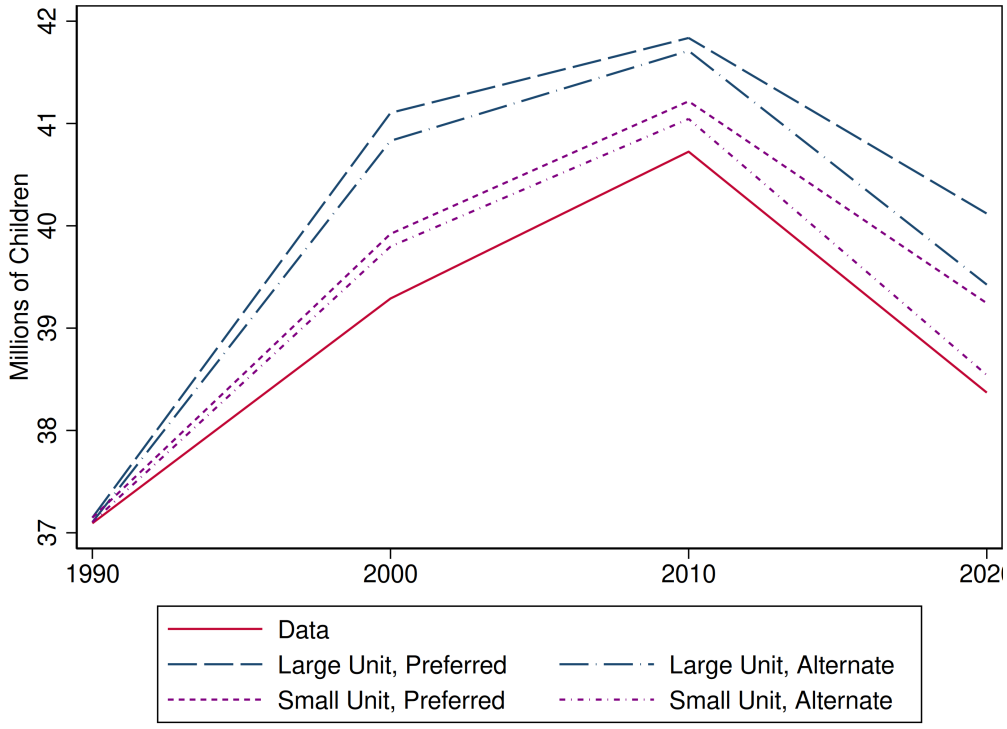


Figure 22: Counterfactual: Children

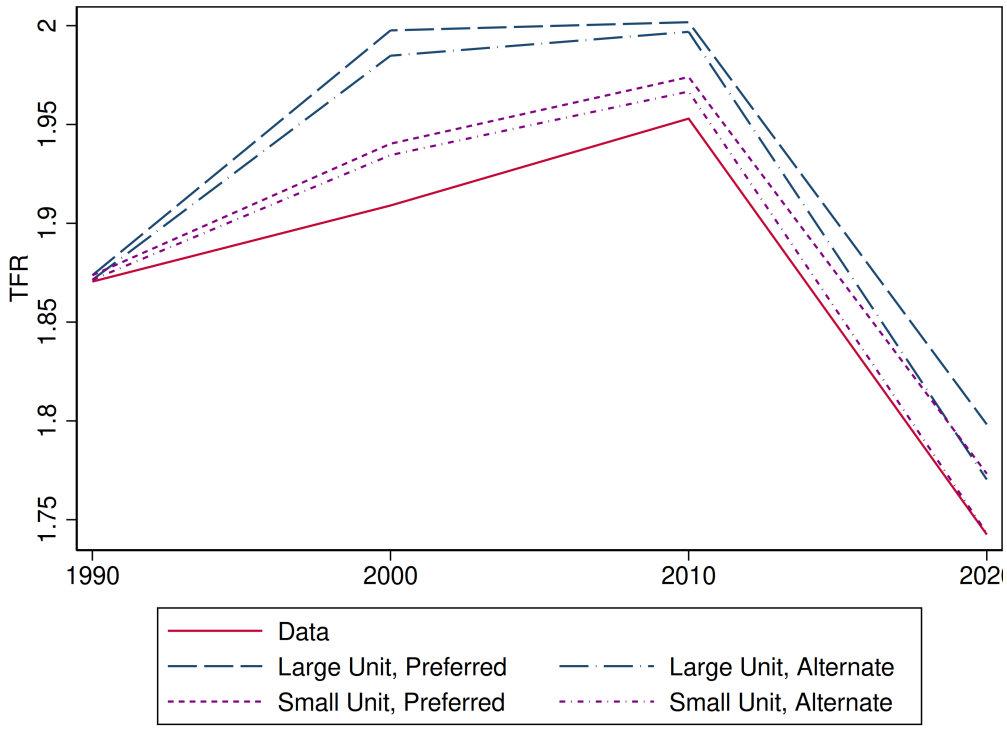


Figure 23: Counterfactual: TFR

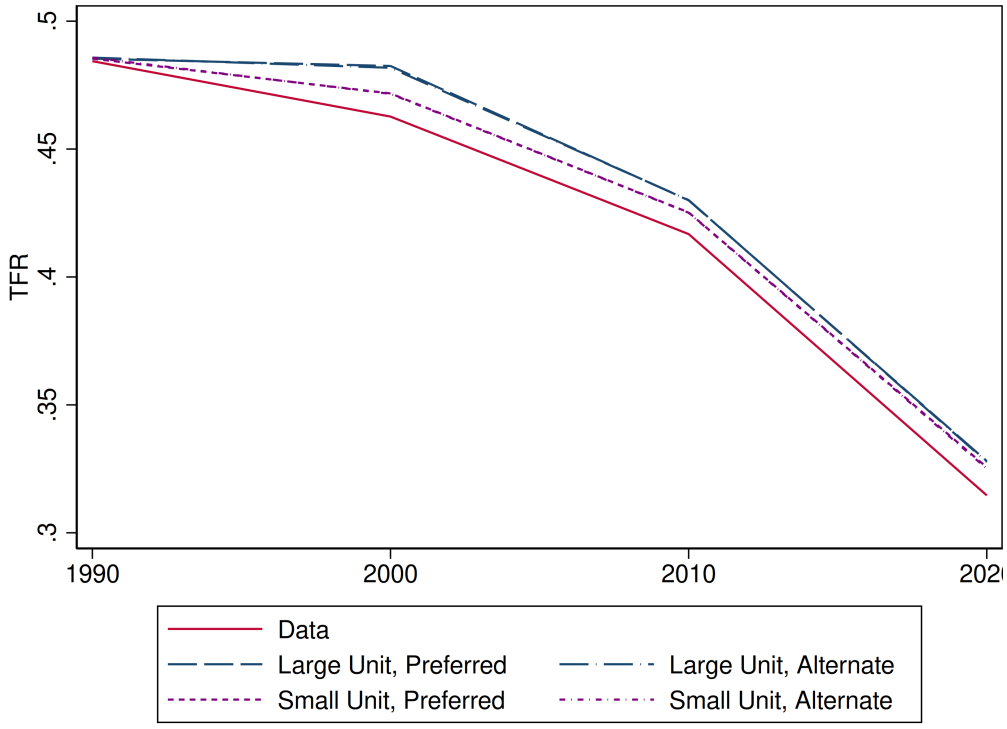


Figure 24: Counterfactual: Share Family, 20-29

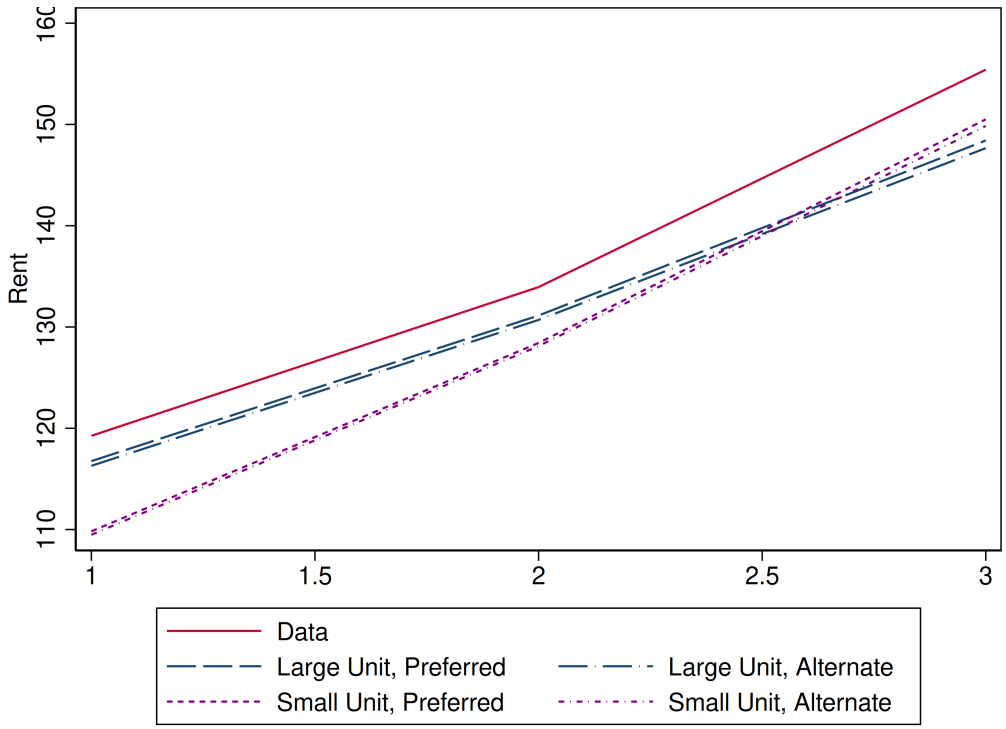


Figure 25: Counterfactual: Rent, 2020