Build, Baby, Build: How Housing Shapes Fertility

Benjamin K. Couillard Job Market Paper

August 25, 2025 PRELIMINARY DRAFT DO NOT DISTRIBUTE

Abstract

Below-replacement fertility leads to demographic crisis. Children increase housing demand, so rising housing costs may be contributing to declining fertility. Most reduced-form papers studying this relationship use geographic variation, and so their results contain some amount of sorting (even under random assignment). To approach this issue, I estimate a dynamic model of joint housing and fertility choice. The housing demand part of the model leverages and extends IO "micro-moment" techniques to estimate a neighborhood choice model with rich demographic heterogeneity without access to the underlying joint distribution, a first in urban economics. A Beckerian quantity-quality model predicts and structural estimates verify that larger families are more cost-sensitive, and so rising housing costs disincentivize fertility. A decomposition exercise shows that if rents had not risen at all since 1990 then there would be 11% more children born and the difference between replacement and 2010s TFR would have shrunk by 90%: this upper bound on housing policy effectiveness show that solving the housing crisis would contribute materially to solving the fertility crisis. Housing supply counterfactuals comparing standard YIMBY versus pro-natal housing policies are in progress.

1 Introduction

Fertility collapse has emerged as a pressing social problem in countries like South Korea (TFR=0.75, 2024), Taiwan (0.89, 2024), Japan (1.15, 2024), Spain (1.12, 2023), Italy (1.18, 2024), and Canada (1.26, 2023), and threatens to do so elsewhere: two thirds of the global population live in a country with below replacement-rate fertility, and the global fertility rate is expected to continue declining in long-run projections through 2100 (UN, 2022). Persistently very low birth rates result in fewer young people and an increasing dependency ratio as the population pyramid inverts. Empirical and theoretical evidence indicates that per-capita pension and healthcare costs increase, a smaller labor force generates inflationary pressure, and reduced innovation and dynamism impede growth (Bloom et al., 2024). With a smaller military-age population, liberal democracies will face challenges in geopolitical competition with authoritarian and totalitarian regimes. Low fertility generates a "demographic dividend" in the short run in the form of fewer dependents, higher per-child parental investment, and greater (mostly female) labor force participation (UN, 2022; Bloom et al., 2024), which may mask the deleterious long-run impacts of fertility decline from policymakers and the public. Although immigration may also temporarily offset these effects, immigrants tend to assimilate in terms of family size (Parrado and Morgan, 2008; Blau, 2015; Andersson, 2004) and virtually all developing countries are currently and are

projected to continue undergoing their own demographic transition (UN, 2022), suggesting that immigration is insufficient as a long-term solution. To avoid the ills of fertility collapse, policy must be used to encourage more births.

Given that housing is generally the largest single household expenditure, and that housing consumption tends to increases with family size[1], it follows that rising housing costs have contributed to declining fertility to some extent, and a reduced-form literature has found supporting evidence in a variety of contexts. However, most papers in this literature use geographic variation in housing costs and fertility, and causal estimates would contain a sorting effect even if housing costs were randomly assigned, as households would sort over locations on the basis of housing costs and their own preferences over family size.

To build on the reduced-form literature, I account for sorting explicitly in a dynamic model of joint housing and fertility choice. Agents make a dynamic living arrangement decision – including fertility – knowing how it will affect their current housing demand (and rationally forecasting the impact on future values). A Beckerian quantity-quality model predicts, and structural estimates verify, that larger families prefer larger and more costly housing units while also being more sensitive to costs. The key mechanism of this paper is that rising housing costs reduce the value of having children to a greater extent than other living arrangement choices, via the greater cost-sensitivity of larger families.

I specify a logit housing demand model where the choices are combinations of neighborhood and unit size, with heterogeneity over household head age, tenure, family / non-family status, and household size, using published marginal distributions of individual dimensions of heterogeneity with choices as sufficient statistics for demand parameters – this is common practice in IO, but the first time it has been exploited in urban economics. I handle endogeneity of rent with the traditional sorting model donut IV where the exogenous (dis)amenity is noise pollution from railways, and contribute a new falsification test for the donut IV strategy. I extend the work of IO economists by showing how the Iterative Proportional Fitting (IPF) algorithm combines all the necessary information in the marginal distributions into a single synthetic joint distribution that generates identical point estimates when used in logit sorting model estimation. Standard IO random-coefficients estimators load demographic variation onto the mean and variance of the coefficients, whereas using a synthetic joint distribution allows utility residuals to vary over demographic groups, which likely has material implications for parameter estimates and allows the study of behavior of specific demographic groups in baseline and counterfactual scenarios with greater fidelity. Finally, I show how certain OLS and 2SLS models also generate identical point estimates regardless of whether the true or synthetic joint distribution is used.

The dynamic living arrangement choice model takes static 10-year housing utilities as given. Agents choose to live alone, with roommates, with parents or other family, to start families of their own and have children, and to convert to a home-owning household, with each option mapping onto an observed heterogeneity housing demand type as discussed above. Although moving costs are assumed to be zero over a 10-year horizon – allowing housing choices to be excluded from the dynamic state – agents account for the likelihood that their adult children will remain with the household, that other family members will choose to live with them, and that their partnership may dissolve, leaving one agent as a single parent and the other as an unattached single. I estimate living arrangement parameters to fit living arrangement data

[1]Indeed, the USDA estimated housing to be the largest child-rearing expenditure at 29% of the total, larger than food, childcare/education, transportation, healthcare, clothing, and miscellaneous (Lino et al, 2015).

exactly, and hold these parameters constant in decomposition and counterfactual exercises in order to isolate the impact of housing on fertility and other living arrangement decisions.

To understand the extent to which rising housing costs have reduced fertility since 1990, I hold rents fixed at their 1990 level in a decomposition exercise. This also serves as a counterfactual exercise where housing supply is perfectly elastic and which provides an upper bound on how maximally accommodative housing policy could have affected fertility. Holding all non-housing determinants of fertility decisions fixed, I find that there would have been 11% more children born and that the 2010s total fertility rate would have been 2.06 compared with 1.74, erasing 90% of the gap with the replacement rate – both figures smaller and more reasonable than that implied by an application of reduced-form estimates. A more traditional counterfactual exercise is in progress, where I compare the fertility impacts of subsidizing the construction of small versus large units. The former is the upshot of standard YIMBY policy, which may stimulate fertility by lowering rents of larger units by redirecting the demand of singles and reducing the likelihood that young adults will live with their parents, while the latter is a more straightforward pro-natalist housing policy.

The paper proceeds as follows. Section 2 reviews the empirical fertility economics literature, discussing how housing costs affect fertility by shifting the budget constraint of the standard model and comparing these effects with other in-kind policies. Section 3 describes the data and provides some descriptive statistics. Section 4 introduces the identifying variation used in structural estimation, verifies the standard reduced-form inverse relationship between rents and fertility. Section 5 shows how the core mechanism of the model – rising prices hurting the value of being a large family more – emerges from a standard quantity-quality model with housing and is verified with an additional reduced-form exercise; it concludes by providing a roadmap to the structural model and its estimation. Section 6 covers iterative proportional fitting and housing demand estimation. Section 7 details the dynamic fertility part of the model and estimation thereof. Section 8 explicates the decomposition exercise, and Section 9 the counterfactual exercise. Section 10 concludes.

2 Literature Review

Writing around the time that the birth control pill gave women true reproductive choice, Becker's (1960) foundational quantity-quality model gave parents preferences over both the number of children (quantity) and their investment in them (quality), such that a greater income elasticity of quality than quantity can generate the well-known inverse relationship between income and fertility. This inverse relationship can also be explained by female wages and opportunity cost of childrearing. However, in recent decades in developed countries, the "pervasive" negative relationship between income and fertility has attenuated or reversed[2], and this pattern largely persists only in developing countries (Doepke et al. 2022).

[2]Possible causes for the disappearance of the initial negative relationship include expansion of public schooling reducing costs of creating child quality, marketization of childcare allowing purchase of childcare for an hourly price that may be less than a working woman's wage, extension of high returns to job experience that delay fertility from college to non-college careers, availability of fertility technology for high-income women who delayed fertility for career reasons, explicitly family-friendly policies like public childcare or expanded parental leave reducing the tradeoff between female labor and childcare, more equitable bargaining positions between men and women in advance of the fertility decision, and shifting social norms with regard to working mothers. See Doepke et al.'s sections sections 4.1, 4.2, 4.3, 4.4, 5.1, 5.2, and 5.3 for more detailed theoretical discussion and empirical support for these mechanisms.

However, although the inverse relationship was pervasive, this does not imply that the causal relationship was ever negative, as Doepke et al.'s review highlights an empirical literature that finds that exogenous increases in male income tends to increase fertility (Fleischer and Rhodes, 1979; Borg, 1989; Heckman and Walker, 1990; Black et al,. 2013; Kearney and Wilson, 2018; and Autor et al., 2019). More recent quasi-experimental papers have provided additional results that exogenous increases in income and/or wealth increase fertility (Cesarini et al., 2023; Tsai et al., 2023; Yonzan et al., 2024). Thus, although there has historically been an inverse relationship between income and fertility, there is also extensive evidence that fertility is a normal good that increases with exogenous increases in income, and it is with this understanding that one must approach the reduced-form literature on housing and fertility.

The reduced-form evidence is consistent with the following logic. Housing shifts a family's budget constraint as a cost or as a form of wealth. An increase in rents will then reduce the fertility of renters by reducing the income available for other goods. An increase in prices or interest rates will reduce the fertility of new buyers in the years after buying, by reducing their wealth or reducing their income via higher mortgage payments. However, an increase in prices will increase the fertility of incumbent homeowners by increasing their wealth.

Dettling and Kearney (2014) find that high metro house prices reduce fertility for renters and increase fertility for owners, likely because prices are highly correlated with rents, and the effect on owners is an average of new and incumbent owners dominated by the latter. Likewise, Clark and Ferrer's (2016) preferred person-fixed-effects estimates find that high metro prices increase fertility for owners and decrease it for renters. Using a state-level price index, Lovenheim and Mumford (2013) find the same pattern of heterogeneity but with a null effect for renters. Pan and Yang (2022) find a positive effect of prices on fertility for homeowners in China, Daysal et al. (2020) find the same in Denmark, and Atalay et al. (2017) find the same in Australia. Liu and Zhang (2024) find that increases in prices induced by speculative spillovers from cities that impose purchase restrictions reduce renter fertility. Averaging over all effects, Clark (2012) finds that women delay fertility in more expensive cities, and Simon and Tamura (2009) find that per-area measures of housing costs reduce fertility of the young as well as completed fertility. Turning to housing finance, Dettling and Kearney (2025) find that the invention of the modern 30-year fixed-rate mortgage in post-war U.S. contributed 10% of the Baby Boom. Comparing families with fixed- and adjustable-rate mortgages, Cumming and Dettling (2023) find that mortgage payment reductions from interest rate cuts lead to increases in fertility. With perhaps the best identification strategy of any of these papers, Fazio et al. (2025) find that winning a housing finance lottery increases fertility.

A common element of many of these papers is the use of a city- or state-level measure of rents or prices rather than that of the family itself, often due to either data limitations or a recognition that individual housing costs are highly endogenous, as families choose both their fertility and their quantity and quality of housing. Local housing costs are chosen by families to a lesser extent and an easier target for identification – but they are still chosen to some extent as families make a neighborhood choice decision. This fact is demonstrated by Heckman and Landersø (2020), whose "finding that sorting is a dynamic process across childhood not only questions the interpretation of earlier work that ignores this phenomenon, it also suggests further analyses of the factors determining how, when, and under which conditions families decide where to live, and how these decisions shape children's lives." As will be discussed in the next section, the reduced-form literature's implicit assumption of zero geographic mobility

(or, infinite moving costs) both biases estimates down and away from zero – regardless of the quality of the identification strategy – and renders estimates inappropriate for interpretation as treatment effects. One of the contributions of this paper is to estimate the relationship between housing costs and fertility while explicitly accounting for mobility.

Bergsvik et al.'s (2021) recent review of the quasi-experimental literature on family policy concludes that in-kind policies like childcare, parental leave, and healthcare are effective in increasing completed fertility, whereas the evidence around fertility subsidies is sparse and mixed. Childcare expansion policies increase fertility at longer time horizons and across parities, although the effect sizes tend to be 1 percentage point or less. When the evidence is restricted to the largest policy changes, maternity leave expansion does increase fertility to a greater extent, with effects persisting in intermediate time horizons and in older age groups, indicating that results would be positive if studies with longer time horizons were available. Health insurance access in the U.S. decreases fertility among the young, suggesting a timing-only effect of avoiding unwanted pregnancies, while access to infertility treatment predictably increases fertility among women approaching the end of reproductive age, suggesting a positive effect on completed fertility of highly-motivated parents. Subsidies generate immediate effects of a magnitude between that of childcare and parental leave (see e.g. Milligan, 2005) but longer-run outcomes are generally not examined. One exception is Parent and Wang (2007), who find a null effect on fertility 15 years after a policy change in Quebec. One the other hand, excluded from that review is Yakovlev and Sorvachev (2025), whose IV-event-study approach find that Russia's Maternity Capital program (worth roughly $10,000) increased the total fertility rate by 21% 15 years after implementation. Thus, although there is extensive evidence that children are a normal good, the quasi-experimental evidence for the effectiveness of fertility subsidies is sparse and mixed (to say nothing of their cost-effectiveness), especially vis-a-vis in-kind policies.

Stepping back, the limited effectiveness of fertility subsidies compared to other income shocks or in-kind family policies may reflect mechanisms described by Easterlin's (1978) hypothesis. Easterlin suggested that fertility depends not simply on absolute income, but on how people's actual living standards compare with the material expectations formed during their own childhood. Thus, the willingness to have children may depend more on perceived economic security or status relative to past or peer norms than on incremental changes in disposable income, especially if subsidies are spread thinly over time and offered to everyone. In contrast, policies that significantly alter lifestyle or perceived security – such as generous parental leave, childcare provision, or lower-cost housing – may more effectively shift decisions about childbearing because their impacts are clearer and more immediately felt, even if cash payments could be put toward replacing lost wages or paying for childcare.

Easterlin's concept helps clarify why housing affordability could be especially effective in influencing fertility decisions. If recent generations grew up in relatively spacious detached suburban homes, but today face increasingly high prices for similar family-friendly housing near employment centers, they may perceive themselves as downwardly mobile and opt to avoid the additional expenditure of starting a family. Building more to lower the cost of family-friendly housing could be a potent policy lever – especially if it is accomplished by correcting a market failure and taking harmful regulations off the book – and there is thus a pressing need for a purposeful examination of the effect of housing reform on fertility that corrects the greatest limitation of the extant reduced-form literature.

3 Data and Descriptives

3.1 Data

Most of the data used in this paper comes from the 1990 Census, 2000 Census, 2008-12 5-year American Community Survey, and 2018-22 5-year American Community Survey. I use a combination of tabulations by census tract (reduced-form and housing demand) and microdata (living arrangement / fertility choice). Each microdata sample is a 5% sample of the US population, with a high degree of continuity over the years, while the census tract tabulations are either short-form Census (100% in 1990 and 2000), population estimates based on the long-form census (a smaller share in 1990 and 2000), or population estimates based on the 5% ACS sample in the 5-year ranges.

I aggregate census tracts minimally so that the 48,000 new census tracts contain all major splits and merges between the 1990 and 2020 vintages.

For housing costs, I use the median rent in a tract for 1, 2, and 3+ bedroom housing units. Because rents are highly correlated with house prices, and Census and ACS house prices are self-reported while rent is a market transaction made monthly, I use rents as housing costs for both renters and owners. This data is originally in the form of household counts by rent bins, and I interpolate the median by estimating a linear CDF. Although each cross-section has rich microgeographic variation, they do not match aggregates well, and so I scale each cross section by a multiplicative factor to match aggregates. I do the same for the general fertility rate (births per woman aged 15-44) used in the reduced-form analysis below, matching aggregates published by the CDC.

The railway noise data used for the donut IV for rents originates from the US Department of Transportation. More detail on railway noise and how the IV is constructed is available below in the reduced-form section, while the housing demand section provides more detail about the marginal distributions of household counts used to construct the left-hand-side variable, and the fertility choice section provides more detail about the microdata sample.

3.2 Descriptives

In developed countries and especially North America, starting a family often coincides with the purchase of a particular housing bundle. This bundle typically is comprised of a single-family detached structure with a yard for children to play in and multiple bedrooms for privacy; neighborhood amenities like low crime, good schools, and other children; and ownership of the unit as opposed to renting. Naturally, a new choice of housing also implies moving to a new residence. Housing and fertility are a joint choice in peoples' lives, and so housing and fertility outcomes are jointly determined.

The strength of this joint determination is captured in the regression that is depicted in Figure 1. Using women aged 26-33 who are household heads or the spouse of a household head, I regress the probability of having given birth in the last year on (1) a set of indicators for the number of bedrooms in the home, (2) an indicator for honeownership, (3) an indicator for whether she resides in a single-family detached home, and (4) an indicator for whether she moved into that home in the last five years, as well as (5) a set of year-city-age fixed effects. Evaluated at the average year-city-age fixed effect, the probability of having given birth in the last year for a woman who has rented a single-bedroom unit for more than five years in a multi-family

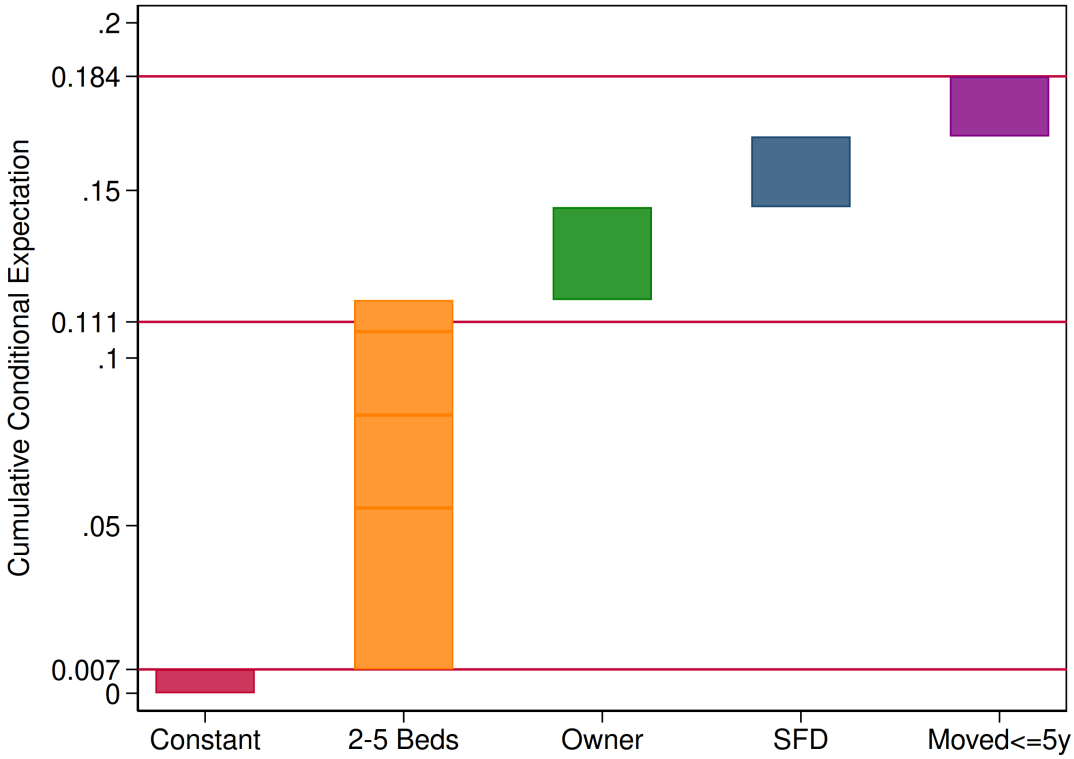


Figure 1

structure is 0.7%. Meanwhile, the probability of having given birth in the last year for a woman who in the last five years moved into an owned 5-bedroom single-family detached home is 18.4%. This increase is 17.7 percentage points and 26-fold increase over the lower probability. The lower probability is 6% of the unconditional mean, whereas the higher probability is 166% of the unconditional mean. Every dimension – homeownership, structure type, duration of residence, and size – individually affects fertility in the expected dimension. From this we can conclude that decisions about housing and fertility are very tightly linked.

Meanwhile, the relative price of this bundle has been increasing for decades, shown in Figure 2. Adjusting for housing quality, resident income, and using 5+ year residents to adjust for changing credit standards, I plot the predicted log price of a 4+ bedroom single-family home and log rent of a single-bedroom 5-19 family home with 1990 levels differenced out. The rent rose about 5% from 1990 to the early 2010s, then a further 15% as of the early 2020s. The price rose nearly 40% from 1990 to 2006, gave up some of these gains as of the early 2010s, then recovered to about the same level as of the early 2020s. As the relative price of the family-friendly housing bundle increases, fewer reproductive-age adults would choose the family-friendly housing bundle, and birthrates would decline.

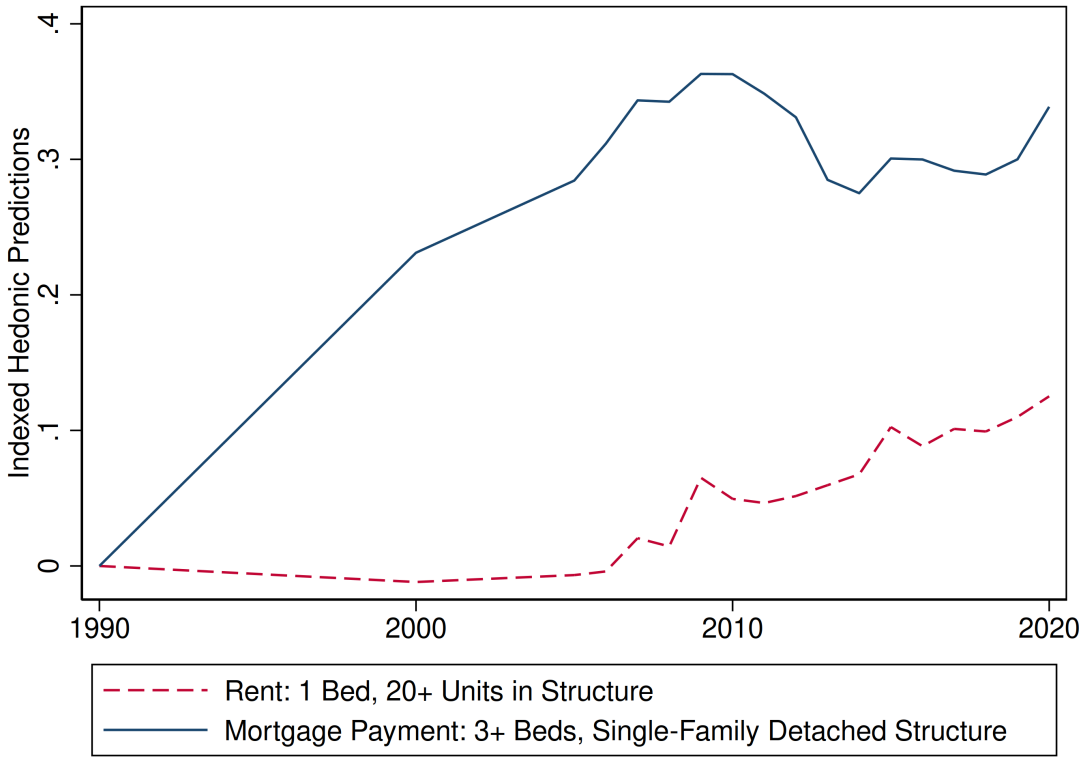

Indexed Hedonic Predictions
0
.1
.2
.3
.4
1990
2000
2010
2020
Rent: 1 Bed, 20+ Units in Structure
Mortgage Payment: 3+ Beds, Single-Family Detached Structure

Figure 2

4 Identification, Reduced-Form, and SUTVA

4.1 Causal Identification With The Donut IV

Causal identification targeting rents in the structural model and reduced-form work in this paper follows from a mechanism first posited by Berry et al. (1995) and adapted to neighborhood choice by Bayer et al. (2007). In the original IO context, price endogeneity arises when firms that produce goods of unobservably (to the econometrician) higher quality mark up their prices more. Under Nash equilibrium for this pricing game, prices are a function not just of other options' prices but also their exogenous characteristics, motivating the use of other options' exogenous characteristics as instrumental variables: the BLP IV.

In sorting models, the logic is distinct. The housing stock is generally fixed in the short run, and potentially very inelastic even in the long run. It generally cannot be produced at a constant cost to meet whatever demand is generated by a given price. In any given time period most housing transactions take place on the secondary market, not on the primary market between buyers and developers at a price that the developer chooses. Every housing unit in some sense competes with every other housing unit in the purchase or rental market. Prices and rents are set by supply and demand, not by strategic interactions between price-setting firms.

However, the logic of the BLP IV can still be used for causal identification in the form of the donut IV, introduced Bayer et al. (2007). The donut IV is the average of some exogenous amenity in the neighborhoods in a donut- or ring-shaped region around a focal neighborhood. Because people often consider broad regions when they are deciding where exactly to live, residence in this region is substitutable with residence in the focal tract. If housing supply is perfectly inelastic in the short run (or merely upward-sloping), an increase in the amenity value of competing tracts (the donut) requires the rent of the focal tract to decline relative to the donut in order to restore equilibrium. Given that amenities are capitalized into prices, and that amenities generally vary continuously over space, the "donut hole" is required to excludes nearby amenities that may be capitalized into the focal tract's price. An outer extent of the donut must be chosen in order to generate a reasonable amount of variation in the IV, and in order to exclude tracts that are not very substitutable with the focal tract, both of which serve to increase the strength of the first stage. The logic of using only the most substitutable competing options to form an IV is very similar to the proposal by Gandhi and Houde (2019) to screen on proximity in product-space to increase the strength of BLP IVs.

The causal mechanism for price-setting that motivates the use of the donut IV suggests a falsification test that does not appear to have been recognized in the literature. When choosing amenities to use as the basis for an IV, the economist generally has a strong intuition about whether the variable constitutes an amenity or disamenity, for all market participants or at least on average. That is, the economist has a strong intuition about whether the variable enters structural utility positively or negatively, and whether it increases or decreases prices. A well-specified donut IV will then be constructed from a variable that enters the price of the focal tract in the expected direction, and the donut IV itself will affect the price of the focal tract in the opposite direction. An amenity in the focal tract should increase prices, and that same amenity in the donut should decrease prices. Likewise, a disamenity in the focal tract should decrease prices, and that same amenity in the donut should increase prices. This is because, given an initial equilibrium, if the amenity value of competing donut tracts increase, the quantity

of housing demanded in the donut will exceed supply and the quantity of housing demanded in the focal tract will fall short of supply. The price of the focal tract must then decrease relative to the donut in order for quantity demanded to increase in the focal tract and decrease in the donut, restoring equilibrium. If an amenity makes a tract more attractive, the donut IV generated from it should reduce prices, and if a disamenity makes a tract less attractive, the donut IV generated from it should increase prices. Although the literature does tend to verify instrument strength with the standard diagnostics, it does not appear to be concerned with the signs of the first-stage coefficients, even though they can indicate that an IV is not well-specified.

The IV used in this paper is constructed from noise pollution generated by rail transportation. The average daily decibel of noise from trains on an extremely fine basis is calculated from data on train counts and speeds as well as noise modeling. This data is generated from data on train counts and speed and noise modeling. It is published by the US Department of Transportation. I use the 2018 data for reduced-form and structural work extending back to 1990, but this should not pose a problem since most rail lines were built long ago. Data is coded as missing if the noise is less than 45 decibels (roughly ambient background noise), which I replace with 40 decibels. I expect noise pollution to negatively impact structural utility and prices, and I validate this later in the paper.

Although the location of rail lines is predetermined, there may be correlated variables that also affect fertility (reduced-form outcome) or structural utility (structural outcome). However, the donut is constructed with a hole around the focal tract which excludes places with amenities that may affect fertility or make a location more desirable to live. Ultimately, whether the identifying assumption is satisfied depends on how amenities that affect fertility or sorting decay over distance, and whether the donut hole is sufficiently large that they have decayed to a negligible level such that the exclusion restriction is satisfied. The next section will discuss the choice of the size of the donut's hole and outer extent. In order for the donut IV's assumptions to be satisfied, rail transportation must also enter structural utility directly as well as through the donut, which may raise issues related to endogenous controls even if the IV is well-specified.

4.2 Rents, Fertility, Endogeneity, And The Instrument

In this section, I discuss and present results from census-tract-level regressions of fertility on rents, instrumenting rents with the same donut IV used for structural estimation. All data save for the IV (discussed just above) originates from the 1990 and 2000 Censuses and 2008-2012 and 2018-2022 5-year American Community Survey samples. To measure fertility, I use the *general fertility rate*: the number of infants (less than one year old) divided by the number of women of reproductive age 15-44. The regressor of interest is the log of the rent of the median 3+ bedroom housing unit in the tract r^3_{jt}. I also use county-year fixed effects $\gamma_{co(j),t}$. I vary the donut inclusion criterion $\mathcal{D}$ over specifications and weight by occupied housing units H_{jt} to calculate it.

$$
\begin{aligned}
y_{jt} &= \gamma_r \ln(r^3_{jt}) + \gamma_x x_j + \gamma_{co(j),t} + \epsilon_{jt} \\
\ln(r^3_{jt}) &= \pi_z z^{\mathcal{D}}_{jt} + \pi_x x_j + \pi_{co(j),t} + \eta_{jt} \\
z^{\mathcal{D}}_{jt} &= \frac{\sum_{k:d(j,k)\in\mathcal{D}} H_{jt} x_j}{\sum_{k:d(j,k)\in\mathcal{D}} H_{jt}}.
\end{aligned}
$$

The endogeneity problem is severe and multidimensional. The intention is to understand the effect of housing costs (rents, highly correlated with prices and mortgage payments) on fertility, but (1) an increase in rents implies an increase in prices that makes incumbent owners wealthier and likely increases their fertility; (2) larger families likely sort to locations where it costs less to provide housing for a larger family; (3) there may be unobserved correlated non-rent amenities that families sort over; (4) there may be unobserved correlated non-rent amenities that impact fertility; and (5) children increase demand for larger units, increasing their rents.

The donut IV alone can address some but not all of these concerns. Because near things are more correlated than distant things, as the donut hole increases in size eventually it will reach a point where the amenities in the donut are uncorrelated with amenities in the focal tract. The IV then plausibly addresses endogeneity concerns (3) and (4) related to unobserved amenities since it is a shifter that is correlated with focal tract rents through the substitution mechanism but not focal tract amenities. As for the reverse causality endogeneity concern (5), a higher local fertility rate may increase local rents, and it may impact local amenities (perhaps through firm entry or greater public good provision to serve a large group of families), but it cannot affect the location of rail lines that were built decades or centuries prior. True, perhaps these places may receive more frequent service and hence greater noise, but there is a distance beyond which this effect would peter out, and so a sufficiently large donut hole would address this concern. However, the donut IV does not address (1) the aggregation of groups (owners and renters) with heterogeneous responses to housing costs in the data[3], nor (2) the selection of people into locations on the basis of their own fundamental preferences over fertility. The structural model will address these concerns.

With many biases in different directions, it is impossible to say *a priori* how the IV estimates will compare with OLS, or how either will compare with the true effect that the structural model is intended to estimate. The composition effect (1) biases estimates up by including a fertility-enhancing wealth effect. Differential sorting on rent (2) biases estimates down as families sort to less expensive places. Differential sorting on correlated non-rent amenities (3) generates an uncertain bias, as does correlated non-rent amenities that affect fertility (4). Reverse causality (5) biases estimates up as the fertility-enhancing effect of lower rents is blunted by the increased housing demand that results from larger families.

4.3 Reduced-Form Results

The first stage coefficients on focal tract and donut rail noise for various donut sizes (each 1km wide) with heteroskedasticity-robust 95% confidence intervals are presented in Figure 3. The smallest donut has a negative coefficient very similar to that on the focal tract, and as the donut moves further away it rises in absolute value, crossing 0 at about 7km and leveling off at positive 0.002-0.003 around 12km until about 35km. The Kleibergen-Paap statistic for these regressions is plotted in Figure 4. For the 10-35km range, it generally exceeds 50 and its highest value is about 90.

The second stage and OLS coefficients on the log rent of the median 3-bedroom unit is plotted in Figure 5. (Only the 10-35km range is plotted due to the failure of first stage diagnostics outside of it.) Except for a spike towards 0 at about 25-30km, the 2SLS coefficients generally lie between -0.04 and -0.06.

[3]The Census Bureau does not publish tables of population counts by single-year age, tenure, and tract.

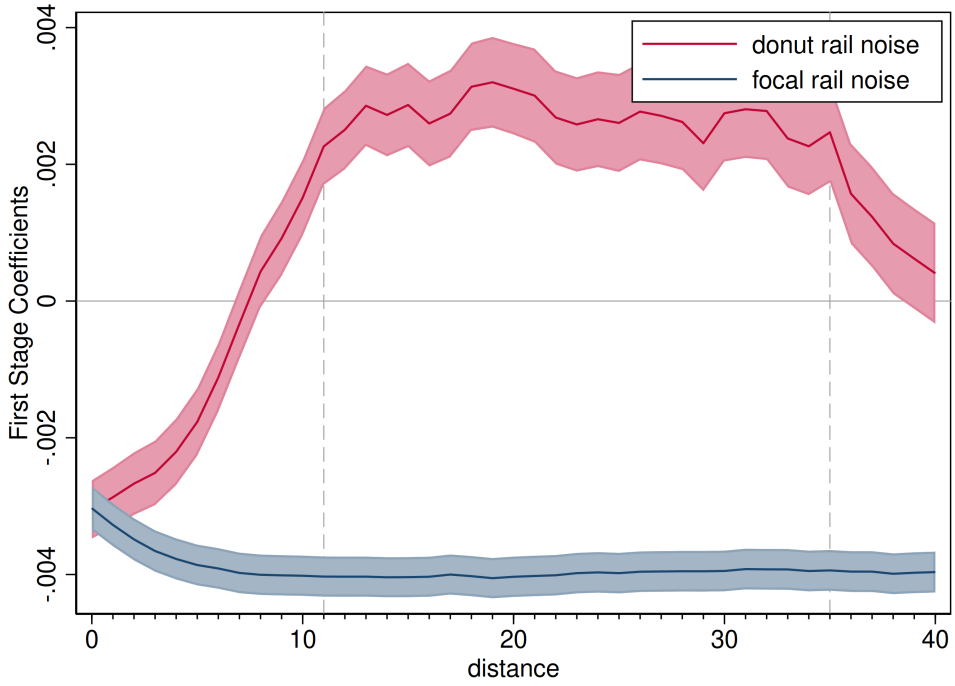


Figure 3

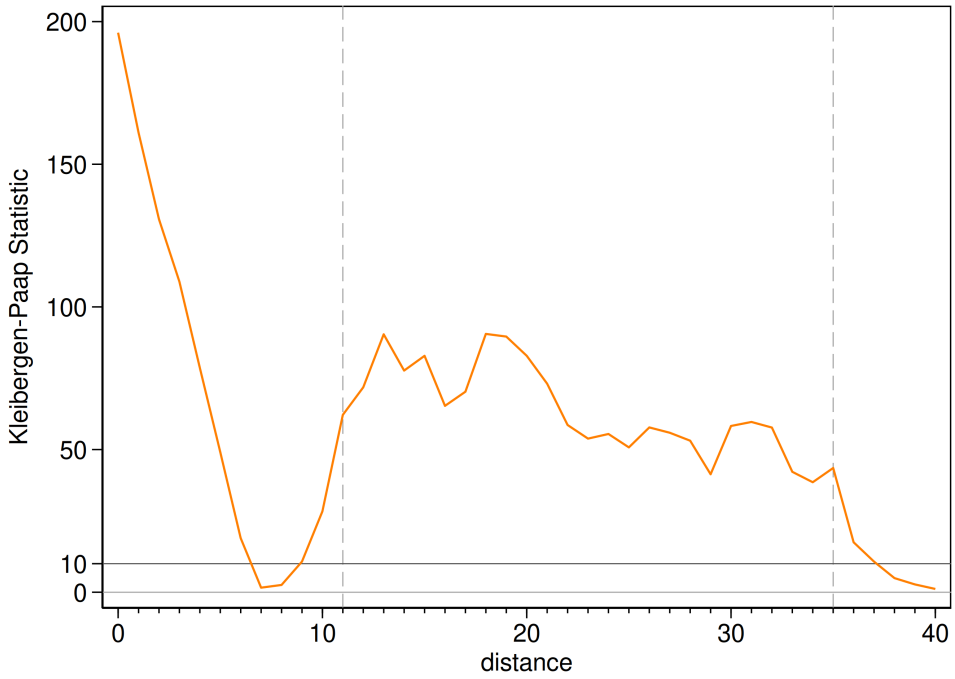


Figure 4

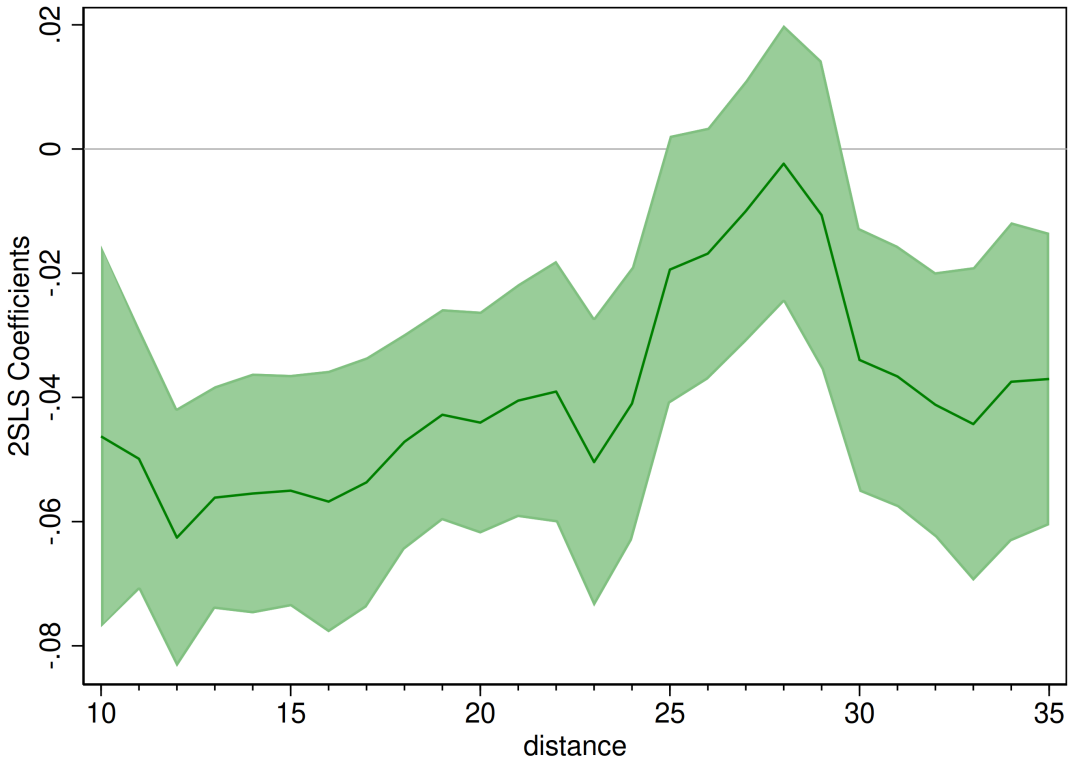


Figure 5

The monotonic rise and sharp leveling-off of the first stage coefficient, as well as the stability of the second stage coefficient within this range, lends credence to the idea that there is a distance beyond which endogeneity is purged (to the extent this is possible with reduced-form analysis), and that this distance is about 11km.

Table 1 presents results for the OLS specification, the first and second stages of the strongest single kilometer donut (13km), and a donut comprising the 11-35km region. Correcting for biases (3), (4), and (5), the 2SLS estimates are more negative than the OLS estimates. Pooling the entire region where the coefficients are stable results in a larger and less precisely-estimated first stage coefficient, resulting in a first stage that is stronger on net and a second stage that is much more precise with a slightly smaller (in absolute value) coefficient. My preferred estimate is this final specification, where the Kleibergen-Paap Statistic is 272 and the rent coefficient is -0.0467.

4.4 SUTVA Violation

The ideal for reduced-form analysis includes both (1) small-area analysis and (2) rich heterogeneity, because (1) to use larger geographies is to introduce attenuation bias in housing costs (since they are highly variable between neighborhoods) and to eliminate potential sources of identifying variation, while (2) to omit heterogeneity is to introduce a SUTVA (Stable Unit Treatment Value Assumption) violation as fertility rates depend on rents in other locations. However, this data

	OLS	First	2SLS	First	2SLS
Ln Med. 3-Bed Rent	-0.0185***		-0.0561***		-0.0467***
	(0.000267)		(0.00915)		(0.00500)
IV: 13km		0.00286***			
		(0.000300)			
IV: 11-35km				0.00871***	
				(0.000528)	
Rail Noise	0.000113***	-0.00403***	-0.0000374	-0.00393***	0.000000317
	(0.0000117)	(0.000149)	(0.0000389)	(0.000149)	(0.0000236)
N	183108	183108	183108	183108	183108
IV			13km		11-35km
KP Stat			90.40		272.3

Standard errors in parentheses

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

Table 1

does not exist for the US, so reduced-form analysis must compromise on some dimension[4].

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:

$$u_{jf}^i = \alpha_f^{\tau(i)} + \beta(1 + f)r_j + \epsilon_{jf}^i.$$

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

$$P_{jf}^\tau = \frac{\exp(\alpha_f^\tau + \beta(1 + f)r_j)}{\sum_k \exp(\alpha_0^\tau + \beta r_k) + \exp(\alpha_1^\tau + 2\beta r_k)}.$$

The fertility rate for a given type and location is

$$F_j^\tau = \frac{P_{j1}^\tau}{P_{j0}^\tau + P_{j1}^\tau} = \frac{\exp(\alpha_1^\tau + 2\beta r_j)}{\exp(\alpha_0^\tau + \beta r_j) + \exp(\alpha_1^\tau + 2\beta r_j)},$$

where notably the identical denominators cancel out of the logit choice probabilities. If these heterogeneous fertility rates are regressed on rents with a valid identification strategy (or, if

[4]A researcher could use small-area age and gender tabulations as I have above, which compromises on heterogeneity. A researcher could use Census and ACS microdata, which includes a great deal of demographics but compromises on geography. A researcher could use confidential natality data, which includes the county the baby was born in (often but not always the county the mother resides in) and some demographics for the mother, slightly compromising both dimensions.

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.

On the other hand, the aggregate fertility rate would be a weighted average of the heterogeneous fertility rates, where the weights are the proportions of either type that choose to settle in j:

$$F_j = W_j^1 \frac{P_{j1}^1}{P_{j0}^1 + P_{j1}^1} + (1 - W_j^1) \frac{P_{j1}^0}{P_{j0}^0 + P_{j1}^0}.$$

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:

$$W_j^1 = \frac{N_j^1}{N_j^0 + N_j^1} = \frac{N\chi(P_{j0}^1 + P_{j1}^1)}{N(1-\chi)(P_{j0}^0 + P_{j1}^0) + N\chi(P_{j0}^1 + P_{j1}^1)},$$

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_j^τ, 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.

5 Theory And Structural Model

To build on the existing reduced-form literature, I specify and estimate a model of joint housing-fertility choice with microgeography and rich heterogeneity, which significantly mitigates the sorting bias discussed in the previous section as well as the composition bias from combining owners and renters. Its structure permits separate estimation of housing and fertility parameters while the heterogeneity allows higher rents to affect the value of having a large family more, disincentivizing fertility. These patterns of heterogeneity are linear approximations to utility that are based on an applied theory model of housing, fertility, child quality, and parental consumption choice, and validated with an additional reduced-form exercise. I discuss that model and reduced-form exercise in this section, before concluding the section with a roadmap for the structural model.

5.1 Applied Theory Model

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

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

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[5]. 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.

Using weak separability[6], we can rewrite the problem as a nested problem[7], 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:

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

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.

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

[6]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.

[7]I prove this in the appendix.

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:

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

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.

$$\begin{aligned} d\ln(nrh) = d\ln n + \cancel{d\ln r} + d\ln h &= d\ln n\Big(1 + \frac{d\ln h}{d\ln n}\Big) > 0 \iff \epsilon_{h,n} > -1 \\ d\ln(ni) = d\ln n + d\ln i &= d\ln n\Big(1 + \frac{d\ln i}{d\ln n}\Big) > 0 \iff \epsilon_{i,n} > -1 \end{aligned}$$

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. Empirical evidence that families with more children consume more housing and are thus more rent-sensitive is also presented in an appendix.

5.2 Roadmap For Structural Model

The purpose of the model is to account for mobility in response to housing costs in a way that the reduced form does not. It does this by specifying a joint choice over many variables that include fertility and location. Fertility is just one part of a dynamic living arrangement choice with sufficient richness to accommodate the breakdown in traditional household structure brought on by the Second Demographic Transition, and the housing decision also includes the size of the house (in bedrooms) as well as tenure.

The model is made tractable by an assumption of (1) additive separability of housing and fertility utility and (2) frictionless geographic mobility. Thus, geography does not enter the state, and the fertility and housing parts of the model can be separately estimated. The assumption of zero moving costs is plausible considering that the time period of the model is a decade, is a more benign corrective to the *de facto* reduced-form assumption of infinite moving costs, and is conservative in the sense of understating the impact of housing costs on fertility since people can easily move to avoid them.

Housing is handled with a standard sorting model, with rich preference heterogeneity coming from IO "micro-moment" conditions hitherto unexploited in urban economics. This heterogeneity is essential for capturing how housing demand changes with the fertility choice, as larger families demand larger houses and are more sensitive to housing costs. Agents make their living arrangement decision understanding how it will affect their valuation of different housing options. When housing costs go up, the value of having a large family declines by more, and so fewer children are born.

Let z index elements of the exogenous state, and x index elements of the endogenous state. Let g index elements of the living arrangement choice, and h index elements of the housing choice. Let ϵ be an EVT1-distributed preference shock on the (g, h) options, and $\mathcal{G}(x)$ be the set of options available to an agent with exogenous state x[8]. Then, the value of a given state to an individual i is

$$V(z, x, \epsilon^i) = \max_{(g,h)\in\mathcal{G}(x)} \{v_{gh}^{zx} + \epsilon_{gh}^i\}.$$

The systematic value of a choice to an agent in a state is v_{gh}^{zx}, and the additive preference shocks give rise to the familiar logit probabilities as a function of these values:

$$P_{gh}^{zx} = \frac{\exp(v_{gh}^{zx})}{\sum_{(g',h')\in\mathcal{G}(x)} \exp(v_{g'h'}^{zx})}.$$

The key assumption is that the value can be decomposed as:

$$v_{gh}^{zx} = u_h^{z,\tau(g,x)} + \alpha_g^{zx} + c_g^{zx}.$$

There is a utility from the housing choice that depends on the agent's state and their living arrangement choice, $u_h^{z,\tau(g,x)}$. There is a utility from the living arrangement choice *per se*, α_g^{zx}. And there is a continuation value from the fact that the living arrangement choice today affects the state in ten years but the housing choice does not, c_g^{zx}.

The payoff of assuming separability and zero moving costs is that the housing and fertility parts of the model can be separately estimated, as housing choice probabilities can be used in housing demand estimation without regard to the living arrangement part of the model[9], and

[8] Available options do not depend on the exogenous state z, and all possible housing options h are available in all states.

[9] To see this, first aggregate over all living arrangement choices to obtain a pure housing choice probability:

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

The numerator of the 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:

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

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:

$$P_h^{z,\tau} = \exp(u_h^{z,\tau}) K^{z,\tau}.$$

We can take $K^{z,\tau}$ to be a logit denominator. If $K^{z,\tau}$ is not equal to the logit denominator, any other part of it is a

the living arrangement part of the model can be estimated using only housing utility and not the housing demand parameters. This is because the separability and zero moving cost assumption allows fertility utility to be factored out of aggregate fertility choice probabilities:

$$P_g^{zx} = \sum_h P_{gh}^{zx} = \frac{\exp(\alpha_g^{zx} + c_g^{zx}) \sum_h \exp(u_h^{z,\tau(g,x)})}{\sum_{(g',h') \in \mathcal{G}(x)} \exp(v_{g'h'}^{zx})}.$$

The sum $\sum_h \exp(u_h^{z,\tau(g,x)})$ is a housing utility aggregator. Conditional on this housing utility aggregator, fertility utility $\alpha_g^{zx} + c_g^{zx}$ is whatever exactly fits aggregate fertility data and thus contains all non-housing determinants of fertility. The estimation procedure to separate these two terms is discussed in detail in Section 7.

The linear parameters of housing utility matter for counterfactual analysis, governing how changes in housing characteristics affect fertility and housing decisions. Estimation of these parameters is discussed in Section 6. With fertility preferences held constant, decreases in rents due to additional housing supply will shift the housing utility aggregator more for larger families due to estimated preference heterogeneity, and so more agents will choose to have large families.

6 Housing Demand

Agents in this model are decision-making adults. The characteristics of the household that they belong to are determined in the living arrangement part of the model and are taken as given in the housing demand model. These households vary by the head's age a, tenure e, family/non-family type f, and size s, and thus housing demand type $\tau = (a, e, f, s)$. The household chooses a neighborhood j and a number of bedrooms $b \in \{1, 2, 3+\}$, with its choice weighted by the number of adult decisionmakers in the household. Non-family households have decisionmakers equal to its size s, while family households assumed to have two decisionmakers.

Housing utility is

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

The explanatory variables are:

- r_{jbt} is the median rent of a unit in neighborhood j of size b at time t
- β_b are fixed effects for size b
- m_{jt} is the share of units that are single-family in neighborhood j at time t, an amenity
- x_j is the noise level from rail transportation, a disamenity that is used to generate the donut IV
- $\phi_{c(j),t}$ is a city-year fixed effect.

Rent r_{jbt} and bedroom fixed effects β_b are the variables of interest. Single-family units are conducive to fertility but the data does not permit adding another dimension to the choice, so instead this concept enters as a share m_{jt} that also helps to capture the general amenity value

constant added to the utility of all choices h given z, τ, leaving choice probabilities unchanged and thus immaterial. Then, housing utilities $u_h^{z,\tau}$ are identified from $P_h^{z,\tau}$ up to an additive constant, just as they would be if there was no fertility part of the model.

of a neighborhood. Rail noise x_j is another amenity that is more plausibly exogenous than m_{jt} and thus more suitable for calculating a donut IV from. Finally, the city-year fixed effect ϕ^{τ}_{ct} captures any time-varying, type-specific amenity at the city level.

The following additively-separable structure is placed on $\beta^{\tau}_r, \beta^{\tau}_b, \beta^{\tau}_m, \beta^{\tau}_x$:

$$\beta^{\tau} = \beta + \beta^a + \beta^e + \beta^f + \beta^s.$$

This structure assists with interpretability and, as will be shown later, is essential for identification when using micro-moments. Also, as will be shown later, this structure need not be and is not placed on $\phi^{\tau}_{c(j),t}$.

6.1 Micro-Moment Identification And Iterative Proportional Fitting

This model can be estimated even without observing the joint distribution of tract, unit size, head age, tenure, family/non-family type, and household size through the use of micro-moments, as has been standard in IO since Petrin (2002) and BLP (2004). The urban economics literature has apparently been slow to recognize that although the joint distribution, necessary for more direct estimation, is observable only in research data centers, a collection of sufficient marginal distributions – micro-moments that are sufficient for estimation – is published by the Census Bureau. I extend the use of marginal distributions for estimating heterogeneity in sorting models to the construction of synthetic joint distributions via the Iterative Proportional Fitting (IPF) algorithm that provide identical point estimates when used in place of the true joint distribution in a BFM-style two-step estimator, by exploiting the sufficiency of marginal distributions in the logit (and Poisson) likelihoods when an assumption of additive separability is placed on preferences. I also provide conditions under which IPF-generated data (and some functions thereof) provides identical point estimates to the true joint distribution when used on the left-hand side of OLS and 2SLS regressions.

These results represent a contribution to (1) empirical IO, (2) urban economics, and (3) applied microeconomics and (4) the myriad fields that already use IPF-generated joint distributions in empirical work. First, joint estimation of discrete choice models with micro-moments, although efficient, does not provide an estimate of the underlying joint distribution, or market shares for specific demographic groups. Estimating the synthetic joint distribution in a preliminary or supplementary step allows the econometrician to study baseline and counterfactual behavior for specific groups. For example, standard joint estimation would not allow me to compare how large families sort geographically under baseline and counterfactual scenarios, unless we discard group-specific utility residuals and use only utility residuals that are homogeneous over demographic groups. Second, the sorting literature in urban economics has not employed estimation with micro-moments of any kind, prematurely ruling out models with both microgeography and rich heterogeneity as the joint distribution is unobserved, even though the US Census Bureau (and likely statistical agencies in other countries) do in fact publish the necessary data. Third, although there is apparently no known way to jointly estimate a linear model with micro-moments when a necessary joint distribution is not observed, the IPF-generated synthetic joint distribution (and some functions thereof) can – depending on the desired model – provide identical estimates when used as an explanatory variable in OLS and 2SLS models, opening up new possibilities in the reduced-form literature. Fourth and finally, these results are also relevant to the fields that use IPF-generated joint distributions, including transportation planning, urban planning, public

health / epidemiology, and political science.

Iterative Proportional Fitting (IPF) fits an estimated joint distribution to observed marginal distributions. Looping through marginal distributions, each joint distribution cell that contributes to a given marginal distribution cell is proportionally adjusted such that the sum of these joint distribution cells equals the total in the marginal distribution, until the updates to the joint distribution fall below a tolerance. The algorithm was first used by Deming and Stephan (1940) to adjust survey weights to match ground-truth marginal distributions. Bishop (1967) showed that the algorithm maximizes a family of likelihoods, including the Poisson likelihood. Results from the literature around the IPF algorithm are helpfully compiled in Bishop et al's (1975) textbook. In addition to being used to calculate survey weights – matching population marginal distributions while preserving interactions in the microdata sample – IPF has been used to construct synthetic joint distributions when multiple marginal distributions are observed but the true joint distribution is desired for analysis but not observed.

The empirical IO literature has also used marginal distributions for estimating choice models when the true joint distribution is desired but not observed. Building on the BLP estimator – which famously uses only market-level aggregates to estimate random-coefficients models with IV – Petrin (2002) adds moment conditions he terms "micro-moments" based on separate marginal distributions of choices with various individual dimensions of demographic heterogeneity, supplementing the unobserved heterogeneity from the random coefficients with observed heterogeneity. Whereas the standard BLP estimator would include only market shares and characteristics of different products in different markets to identify mean utility parameters and the spread parameters of the random coefficients, Petrin adds to the GMM objective function the probability that a household will purchase a vehicle conditional on income, the average family size conditional on purchasing different types of vehicles, and the probability that the household head is aged 30-60 conditional on purchasing different types of vehicles, in order to identify utility parameters that are heterogeneous over income, family size, and age. Berry et al (2004) provides additional formalization for estimating this kind of observed heterogeneity using marginal distributions, and a systematic treatment in the PyBLP package is provided by Conlon and Gortmaker (2024).

But although this joint estimation is efficient, it can only produce estimates of the joint distribution that constrain the utility residual ξ_j to be equal over demographics because it loads all demographic variation onto the mean and variance of the random coefficients. This is quite reasonable considering that we do not observe the demographic-varying market shares that would be used to generate a demographic-varying residual. My contribution to empirical IO is to generate an estimate of these demographic-varying market shares using all of the information available in the marginal distributions, so as to estimate a demographic-varying utility residual ξ_j^τ which can be held constant in counterfactuals, and thereby study the difference between factual and counterfactual behavior for specific groups with greater fidelity.[10]

The condition that is required for these sufficiency results is that, for every additive term (except for the error term) in the Poisson estimating equation implied by the preference specification, some marginal distribution with the same indexes must be observed (or derived from

[10]Although I do not include random coefficients because the residential choice context is not very suitable, an IPF-generated estimate of the joint distribution could be used in a random-coefficients model that allows the utility residual to vary over demographic groups, which would likely have material implications for parameter estimates by allowing the utility residual ξ_j^τ to vary over demographics and not loading all demographic variation onto the mean and variance of the random coefficients.

another marginal distribution via aggregation) and entered into the IPF algorithm.[11] Then, true joint distributions and IPF-estimated synthetic joint distributions provide identical estimates when used in choice models with observed heterogeneity. This condition is feasible under the standard assumption of additive separability in heterogeneous preferences (whereas if every fine demographic group were allowed to have its own unstructured parameters then the joint distribution would need to be observed). Then, the marginal distributions – counts of agents choosing each option with individual dimensions of heterogeneity and the distribution of these dimensions of heterogeneity together without choices – are sufficient statistics for preferences in the logit and Poisson likelihoods. When using a BFM-style two-step IV estimator to confront endogeneity, the marginal distributions are sufficient for product fixed effects and heterogeneity parameters in a maximum-likelihood first step, and so the true and synthetic joint distributions will provide identical product fixed effects and thus identical 2SLS baseline preferences in the second step. While estimators using only the marginal distributions can and have been used (in the form of IO micro-moments), the IPF-generated synthetic joint distribution is useful because it enables the estimation of group-specific residuals ξ_j^τ, permits the extension to linear models, and encodes the multiple marginal distributions in a single convenient tabular form. It is my hope that this work will represent a unified framework for understanding what information linear and exponential conditional mean models demand from the data, how parametric assumptions shape the demanded information, how often these information demands can be satisfied using published data when superficially it appears otherwise, and how such information can be combined into a useful form using an elegant and well-understood algorithm.

Owing to the difficulty of analyzing tables with an arbitrary number of dimensions each with an arbitrary number of levels, I prove these results in a three-dimensional context (with a single choice dimension and two demographic dimensions) where all three two-dimensional marginal distributions are observed. This is the simplest relevant context: if there were a single demographic dimension then its distribution with the choice variable would be sufficient for estimation without IPF. One-dimensional marginal distributions are not useful.

6.2 Example

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

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

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.

[11] In the sample discussed in the following section, this means that in order to estimate a sorting model where agent counts depend on α^{as}, $\beta^a x_j$, $\beta^s x_j$, and an error term, we need only observe marginal distribution counts N^{as}, N_j^a, and N_j^s, whereas if we did *not* impose the additive separability assumption $\beta^{as} = \beta^a + \beta^s$ there would be a $\beta^{as} x_j$ term that would then require observing the joint distribution N_j^{as}.

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

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

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

$$N_j^{as} = P_j^{as} N^{as} = \exp(\alpha^{as} + \beta^a x_j + \beta^s x_j + \xi_j^{as}) \quad \text{Poisson Estimating Equation}$$

where

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

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

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

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_j^a, N_j^s are observed while N_j^{as} is not. Note that additive separability on preferences, $\beta^{as} = \beta^a + \beta^s$, generates this condition. If we let $u_j^{as} = \beta^{as} x_j + \xi_j^{as}$ then we would need to observe N_j^{as}.

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

$$N_j^{as} = \exp(\mu_j^s + \mu_j^a + \mu^{as} + \eta_j^{as}).$$

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

$$\hat{N}_j^{as} = \text{IPF}(N_j^a, N_j^s, N^{as}).$$

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

$$\hat{\omega}_j^{as} = N_j^{as} - \hat{N}_j^{as}.$$

Because IPF matches the observed margins exactly, we have

$$\sum_j \hat{N}_j^{as} = N^{as}, \quad \sum_a \hat{N}_j^{as} = N_j^s, \quad \sum_s \hat{N}_j^{as} = N_j^a$$

and so

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

Theorem 1: For any exponential conditional mean model with log-mean additively separable across the $(j,a),(j,s),(a,s)$ margins, $\hat{N}_j^{as}$ provides the same point estimates as N_j^{as} when used as a dependent variable.

Proof: Appendix. The gradients are identical regardless of whether N_j^{as} or $\hat{N}_j^{as}$ is used because they depend only on the marginal distributions, and IPF fits the observed margins exactly. Essentially an application of the well-known sufficiency of marginal totals for fixed effects (Bishop et al (1975), Ch. 3).

Remark: Because models with variables that only vary at the j, s, or a level also lie within this space, this result also applies to such models, such as the first step of the two-step estimator where product fixed effects and heterogeneous slopes are estimated:

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

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.

Proof: Follows directly from identical estimates of δ_j in the first step.

Lemma 1: The maximizer $\hat{\mu}$ for the Poisson log-likelihood

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

satisfies

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

Proof: appendix. Originally proven by Bishop (1967).

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.

Proof: appendix. Follows directly from Poisson score equation for the log-likelihood in Lemma 1 – error is orthogonal to design matrix, hence any variable collinear with design matrix.

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.

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

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

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)

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

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 $\hat{\beta}^N = \hat{\beta}^{\hat{N}}$.

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.

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.

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

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.

Proof: 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.

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

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

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

In summation notation,

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

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:

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

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:

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

6.3 Data and Estimation

Recall the utility specification for housing demand:

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

where β^τ is additively separable in slopes for each level of each dimension of heterogeneity. To estimate this housing demand model, I first estimate the unobserved joint distribution of households:

$$\hat{N}_{0jbt}^{aefs} = \text{IPF}(N_{jt}^{fs}, N_{jt}^{es}, N_{jt}^{af}, N_{jbt}^{e}, N_{p_t(j),b,t}^{aefs})$$

where $p_t(j)$ is 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}.

Then, I obtain the joint distribution of decision-making adults under the assumption[12] that non-family households are all adults and family households have two:

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

This approach is consistent with (1) Census Bureau definitions such that children are never observed in non-family households, and (2) a living arrangement modeling approach, itself required due to data constraints, that adult children and non-nuclear family members and co-residing friends receive a housing utility normalized to zero. Housing decisions for families are assumed to be made by a head and spouse.

An additional step is required before the utility specification conforms to the estimated joint distribution. To see this, expand the rent term, the sole term that varies over both j and b:

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

Following the principle established in the previous subsection, we would need to enter N_{jbt}, N_{jbt}^a, N_{jbt}^e, N_{jbt}^f, N_{jbt}^s, N_{ct}^{aefs} into the IPF algorithm. For the city-level count N_{ct}^{aefs} (necessary to estimate ϕ_{ct}^{aefs}) the presence of distributions of everything save j at the PUMA level (below city)

[12] Because some input marginal distributions do not vary at the (f, s) level ($N_{jt}^{es}, N_{jt}^{af}, N_{jbt}^{e}$), this re-weighting cannot be applied to the input marginal distributions. However, because the IPF-estimated joint distribution of household counts is the best estimate of the joint distribution of household counts (by maximizing the likelihood), applying the weighting after estimation is better than using any guess at re-weighted marginal distributions in the IPF algorithm.

provides more information that is needed, but notably the only tract-level marginal distribution that includes bedroom count available from the Census Bureau is N^e_{jbt}.

The situation can be rectified by discarding some of the variation in r:

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

Then,

$$\begin{aligned}\beta^\tau_r \hat{r}_{jbt} = \beta_r(\hat{\gamma}_{jt} + \hat{\gamma}_{p_t(j),b,t}) + \beta^a_r(\hat{\gamma}_{jt} + \hat{\gamma}_{p_t(j),b,t}) + \beta^e_r(\hat{\gamma}_{jt} + \hat{\gamma}_{p_t(j),b,t}) \\ + \beta^f_r(\hat{\gamma}_{jt} + \hat{\gamma}_{p_t(j),b,t}) + \beta^s_r(\hat{\gamma}_{jt} + \hat{\gamma}_{p_t(j),b,t}),\end{aligned}$$

and now $N^{fs}_{jt}, N^{es}_{jt}, N^{af}_{jt}, N^e_{jbt}, N^{aefs}_{p_t(j),b,t}$ are sufficient[13].

After IPF estimation of the left hand side variable, I estimate heterogeneity parameters with (Poisson) maximum likelihood using the following estimating equation:

$$\begin{aligned}\hat{N}^\tau_{jbt} &= H_{jt} \exp\left(\beta^\tau_r \hat{r}_{jbt} + \beta^\tau_b + \beta^\tau_m m_{jt} + \beta^\tau_x x_j + \phi^\tau_{c(j),t} + \xi^\tau_{jbt}\right) - \hat{\omega}^\tau_{jbt} \\ &= \exp\Big(\delta_{jbt} + (\beta^\tau_r - \beta_r)\hat{r}_{jbt} + (\beta^\tau_b - \beta_b) + (\beta^\tau_m - \beta_m)m_{jt} \\ &\qquad + (\beta^\tau_x - \beta_x)x_j + (\phi^\tau_{c(j),t} - \phi_{c(j),t}) + (\xi^\tau_{jbt} - \xi_{jbt})\Big) - \hat{\omega}^\tau_{jbt}.\end{aligned}$$

Because census tracts vary in size and population (especially after the tract consolidation algorithm), it is necessary to adjust for size in some way using an exposure variable. For this, I choose the total housing stock[14] H_{jt} as the Poisson exposure variable. Then, I collect preferences of the baseline group in δ_{jbt}, which also absorbs the exposure variable, so that the linear parameters capture preferences relative to the baseline group. Due to computational constraints, I randomly divide the data into two subsets 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. I conservatively adjust standard errors for the linear heterogeneity parameters by a factor of $\sqrt{2}$.

I estimate baseline coefficients with 2SLS using a donut instrument over rail noise in a

[13]This data would be sufficient if $\hat{\gamma}_{jt}$ and $\hat{\gamma}_{p_t(j),b,t}$ entered as distinct regressors with distinct coefficients, and to constrain the coefficients to be equal uses strictly less information.

[14]Arguably, it would be better to use H_{jbt}, breaking out the housing stock by unit size. The current approach arguably treats unit size as something that can be adjusted by agents frictionlessly aside from variation in rent, which is unrealistic. However, because agents in this model are adults and not households, either approach treats the presence of crowding as evidence of an option being particularly desirable, and because this happens more in high-amenity expensive locations it can exacerbate endogeneity bias in rent. The problem is less severe when using H_{jt} as the exposure variable. In models where agents are households and when using data from the US Census Bureau where households are equivalent to occupied housing units, H_{jt} is the same as the left hand side variable, and so measures of vacancy would be required for identification. Unfortunately, counts of vacant housing units by tract and number of bedrooms are not available from the Census Bureau, and so that approach cannot be used here.

specified donut interval $\mathcal{D}$, adjusting for exposure H_{jt}:

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

Because the objects considered are similar but distinct to those discussed in the reduced-forms section, it is informative to briefly review the differences between the structural housing demand and reduced-form specifications aside from the obvious difference in outcome variable. Here, there is an additional dimension of variation, unit size, whereas the reduced-form outcome variable could not be observed by unit size. Furthermore, here I include unit size fixed effects. Here, I use the median rent of 1, 2, and 3+ bedroom units, whereas in the reduced-form section I considered only the rent of 3+ bedroom units as these are the most desirable for families. Here, I let rent enter linearly, to better capture the variation in the curvature of the *a prior* homogeneous utility function resulting from the fertility decision, whereas the reduced-form specification used the log of rent. Here, I use Core-Based Statistical Area (CBSA, successor to MSA) fixed effects rather than county fixed effects, since county fixed effects would not be identified by IPF-generated data and PUMAs lack an inherent geographical concept[15].

It is likewise informative to consider the differences in the nature of the endogeneity between the reduced-form and structural specification, which directly addresses the two main concerns about the reduced-form. Because this housing demand model includes heterogeneity over tenure, the straightforward response of renters to higher housing costs is not conflated with the mixed housing cost / capital gains effect of homeowners, which manifests as a less steeply-sloped demand curve. Likewise, this sorting model handles rent-based sorting of heterogeneous groups directly.

More generally, in discrete-choice problems we are fundamentally concerned with firms setting higher markups for unobservably higher-quality goods, or in this case unobserved neighborhood amenities ξ_{jbt} being capitalized into higher prices and rents r_{jbt}. In a more standard empirical IO random-coefficients model, the shocks on the coefficients are assumed to be *i.i.d.* and hence uncorrelated with product characteristics, thus preference heterogeneity can be estimated without IV. Here we make a similar assumption that the relative error terms of different types $\xi_{jbt}^{\tau} - \xi_{jbt}$ is uncorrelated with neighborhood-bedroom characteristics, permitting preference heterogeneity to be estimated without IV. Rather, the IV corrects for correlation between unobserved amenities valued by all types and neighborhood-bedroom characteristics.

In this context, we may be concerned that some neighborhoods may have amenities that are valued more by families or which increase local fertility, and that these neighborhood will have higher large-unit rents compared to small-unit rents, thus invalidating the assumption placed on the heterogeneity-ML stage. This concept is somewhat proxied by the share of housing units that are single-family. If such a violation did occur, it would serve to increase large families' rent coefficient upwards towards zero, and towards the rent coefficients of households without children – a bias that serves to shrink the magnitude of counterfactual results, which depends

[15]PUMAs are created by the Census Bureau in order to meet a population threshold deemed acceptable for balancing data and research value with disclosure risk. As a result, they comprise large geographic regions in rural areas and neighborhoods in urban areas. For example, in 1990 the state of Wyoming had 4 PUMAs, whereas the five boroughs of New York City had about 35.

critically on the extent to which fertility impacts preference heterogeneity. But we may also point towards the possibility of some neighborhoods having amenities that are valued more by singles or which decrease local fertility, which increase the rent of small units relative to large units, and which serve to increase singles' rent coefficient upwards towards zero, and away from the rent coefficients of families. Although such a bias would serve to enlarge the magnitude of counterfactual results, it is *a priori* unclear which bias would dominate if any. Thus for now the assumption appears benign, although in principle it could be avoided by using an estimator that permits joint estimation, such as Mullahy's (1997) IV estimator for count data.

6.4 Results

First stage and baseline-group 2SLS results are shown in Figures 6, 7, 8, and 9. Because they apply to one out of 88 distinct demographic groups, the baseline 2SLS coefficients are meaningful only insofar as they provide an anchor for the preferences of all groups when combined with the heterogeneity results.

Like in the reduced-form, the donut coefficient levels off at about 11-12km and remains stable until about 35km, but this time it does not completely fall to insignificance until nearly 60km, much more distant than 40km in the reduced-form. This provides more evidence for the idea that the characteristics of nearby neighborhoods do enter utility, but that widening the donut hole to exclude them does result in the instrument satisfying the exclusion restriction. The Kleibergen-Paap Statistic reaches about 100 at its highest, similar to 90 in the reduced-form.

Curiously, the baseline rent coefficient falls precipitously past a distance about 25km. Generally, the greater the magnitude of the rent coefficient, the greater is the difference between 2 bedroom and 3+ bedroom units. The more sensitive to rent agents are, the larger must be the value of larger and more expensive units in order to fit the observed choices.

I present OLS and selected 2SLS results (and first stages) in Table 2. The OLS rent coefficient is negative for the baseline group, but not negative enough to ensure that all groups have downward-sloped demand. The 2SLS coefficients are larger by an order of magnitude. Collecting the region where the donut first stage coefficient is stable into a single instrument again results in a larger first stage coefficient and thus much stronger first stage but more shallowly-sloped demand and a smaller difference between the 2 bedroom and 3+ bedroom fixed effects. The 14km donut is the one with the single steepest demand curve and largest difference between the 2 bedroom and 3+ bedroom fixed effects. It is not quite enough to ensure that all groups prefer 3+ bedroom units over 2 bedroom units at equal prices, but it is the closest, and it is thus the least implausible estimate and my preferred specification.

Histograms for parameters of interest are shown in Figure 10. These histograms are weighted by type frequency across all years. Every group has a downward-sloping demand curve, and prefers a 2 bedroom unit to a 1 bedroom unit. Almost every group prefers a 3 bedroom unit to a 2 bedroom unit. The average rent coefficient of -0.0016 for owners becomes -0.0329 (-0.0330 unweighted) when converted to an annual basis and from 2022 to 2000 USD, which is similar to the value of -0.0317 estimated for owners by Bayer et al (2016) using a completely different method[16].

[16]Bayer et al (2016) identified the marginal utility of the numeraire using the response of household moving decisions to realtor fees. Using data over 1994-2004, they estimated that this quantity is $-0.035 + 0.00008\times$income in 1,000s of 2000 USD, which I evaluate at 40.696, median household income at the midpoint of their sample in 1999, to obtain -0.0317. Because their analysis is concerned with homeowners, I compare it with the weighted average of the numeraire

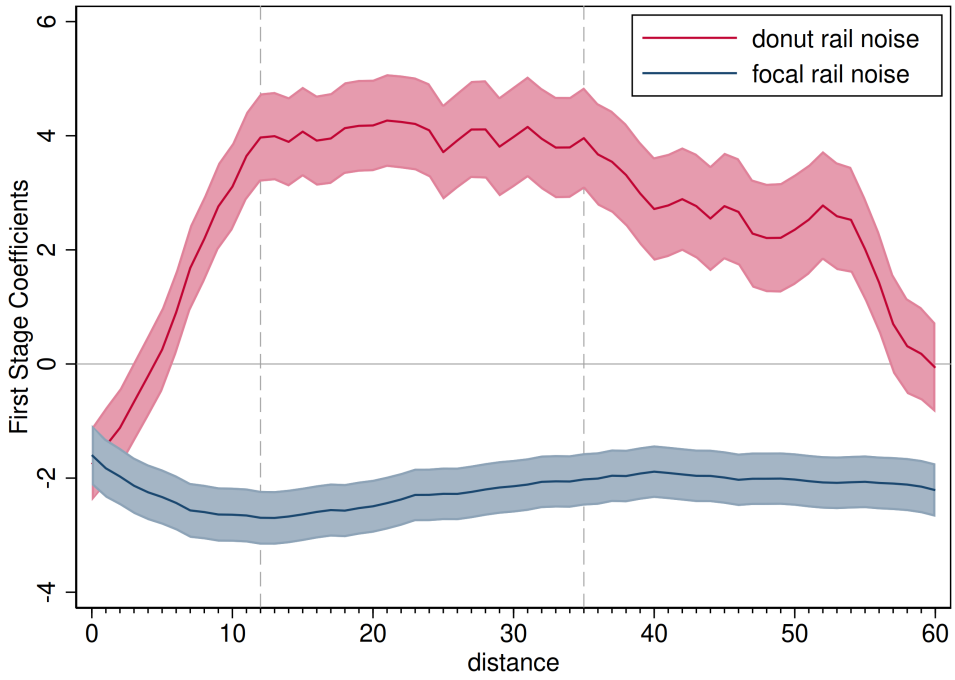


Figure 6

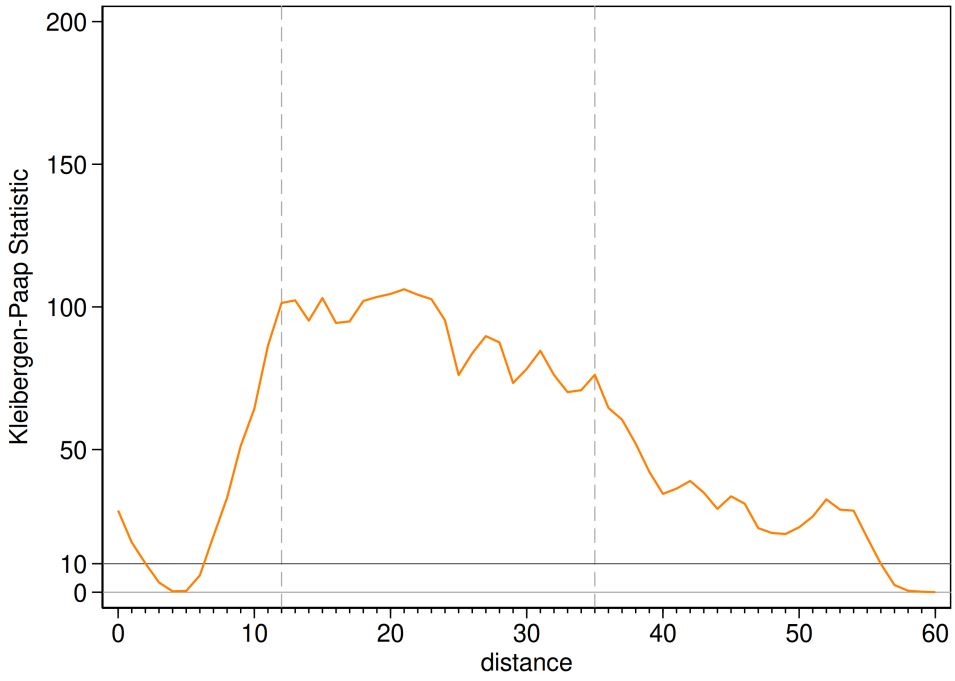


Figure 7

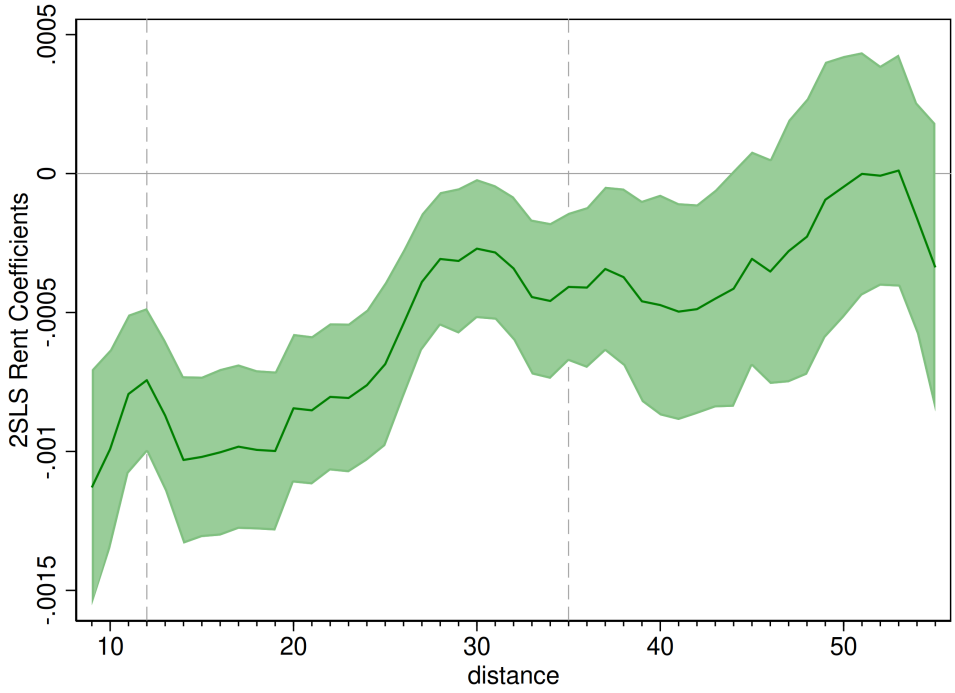


Figure 8

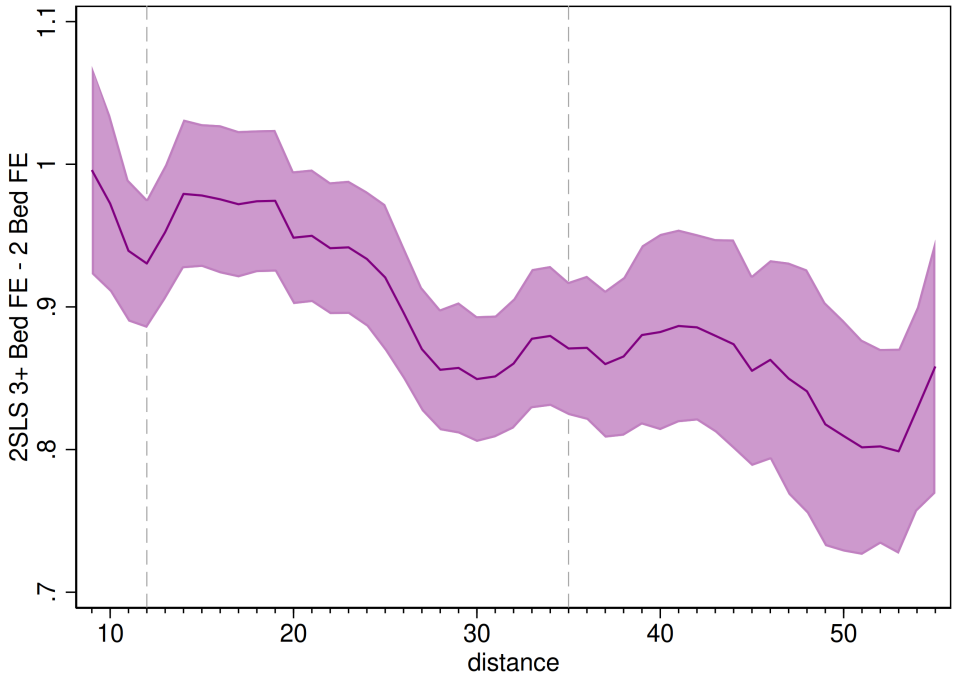


Figure 9

	OLS	First	2SLS	First	2SLS
Rent	-0.0000162**		-0.000679***		-0.00103***
	(0.00000553)		(0.0000714)		(0.000154)
2 Bedrooms	1.146***	159.0***	1.252***	159.0***	1.308***
	(0.00405)	(0.283)	(0.0120)	(0.283)	(0.0247)
3+ Bedrooms	1.952***	330.1***	2.171***	330.1***	2.287***
	(0.00670)	(0.396)	(0.0245)	(0.396)	(0.0511)
Rail Noise, Focal	0.00269***	-2.795***	0.00116***	-2.674***	0.000354
	(0.000255)	(0.237)	(0.000327)	(0.239)	(0.000482)
Rail Noise, 12-35km		9.981***			
		(0.556)			
Rail Noise, 14km				3.896***	
				(0.399)	
N	434370	434370	434370	434370	434370
IV			12-35km		18km
KP Stat			322.4		95.33

Standard errors in parentheses

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

Table 2

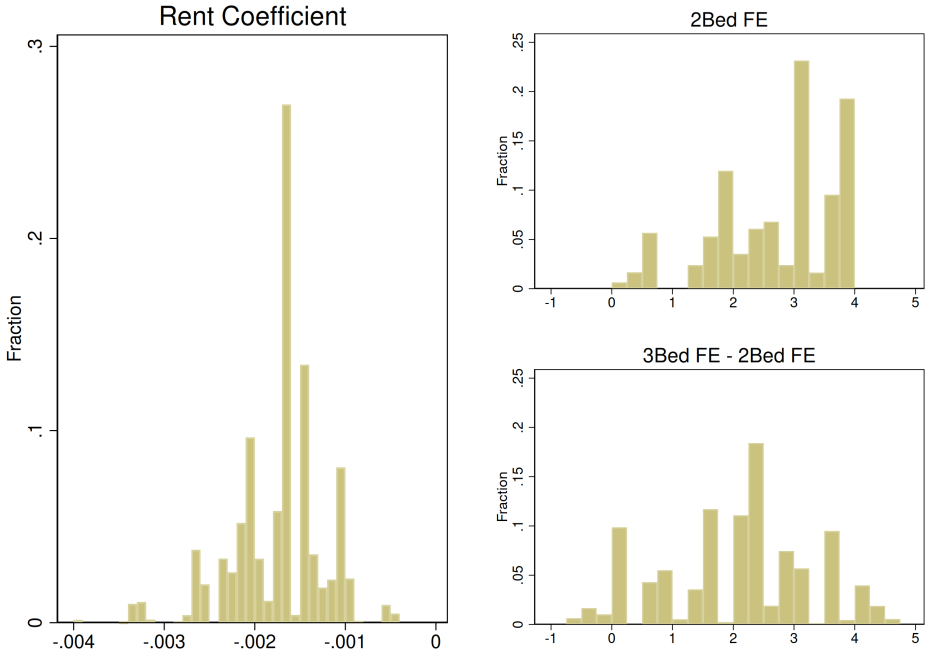


Figure 10

Next, I depict rent coefficients with 95% confidence intervals[17] for renters aged 35-44 by family/non-family and household size in Figure 11. Families and larger families are more rent-sensitive than non-families and smaller families, exactly as predicted by theory. The difference between families and non-families is statistically significant, as are many of the differences between the sizes of each.

The additive separability assumption on the coefficients is clearly visible as the family coefficients are shifted down from the non-family coefficients by a constant, and each curve follows the same shape: heterogeneity over household size is identified by both groups. This explains why the rent coefficient rises from one-person to two-person households: most two-person households are either couples or adult roommates, likely including two earners rather than one earner and one dependent, and thus higher income and less sensitive to rent. However, more three-person households are families rather than unrelated adult roommates and likely to include a dependent, as with four-person, *etc.*, as the demand curves then become more steeply sloped.

Next are the bedroom fixed effects and 95% confidence intervals[18] for renters aged 35-44 by family/non-family and household size in Figure 12. For all groups, the 1 bedroom fixed effect is normalized to zero. Most groups have statistically indistinct 2 bedroom fixed effects,

coefficient for all owners after annualizing rent and converting to 2000 USD, $-0.00156 \times 12 \times 1.76 = -0.0329$.

[17]The 95% confidence intervals include the $\sqrt{2}$ adjustment and are suitable for testing differences between the rent coefficients of different groups. The confidence intervals do not factor in uncertainty from baseline coefficient estimation and are thus not suitable for testing equality to specific values, although we can safely conclude that they are distinct from zero.

[18]The same caveat as for the rent coefficients applies.

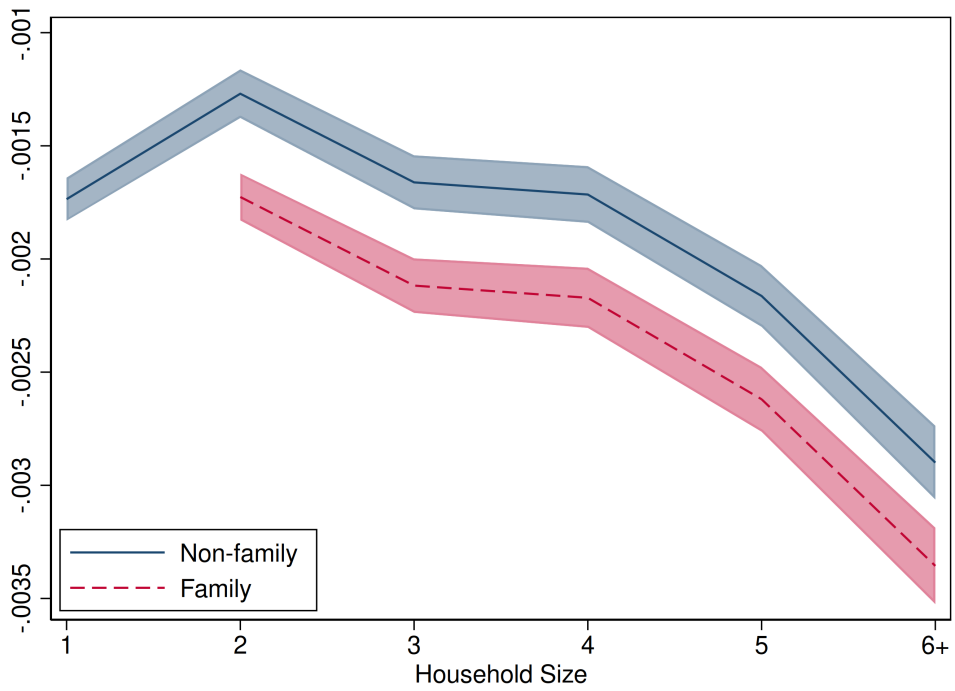


Figure 11

although we can conclude that smaller households have a smaller fixed effect (and hence a smaller marginal utility of a 2 bedroom unit relative to a 1 bedroom unit). However, we can conclude that larger households have significantly higher marginal utility of a 3+ bedroom unit relative to a 2 bedroom unit, for every level of household size.

7 A Dynamic Model Of Fertility Choice

7.1 Setup

Let z index elements of the exogenous state, and x index elements of the endogenous state. Let g index elements of the living arrangement choice, and h index elements of the housing choice. Let ϵ be an iid EVT1-distributed preference shock on the (g, h) options, and $\mathcal{G}(x)$ be the set of options available to an agent with exogenous state x[19]. Then, the value of a given state to an individual i is

$$V(z, x, \epsilon^i) = \max_{g \in \mathcal{G}(x), h} \{v_{gh}^{zx} + \epsilon_{gh}^i\}.$$

The systematic value of a choice to an agent in a state is v_{gh}^{zx}, and the additive preference shocks give rise to the familiar logit probabilities as a function of these values:

$$P_{gh}^{zx} = \frac{\exp(v_{gh}^{zx})}{\sum_{g' \in \mathcal{G}(x)} \sum_h \exp(v_{g'h'}^{zx})}.$$

[19]Available options do not depend on the exogenous state z, and all possible housing options h are available in all states.

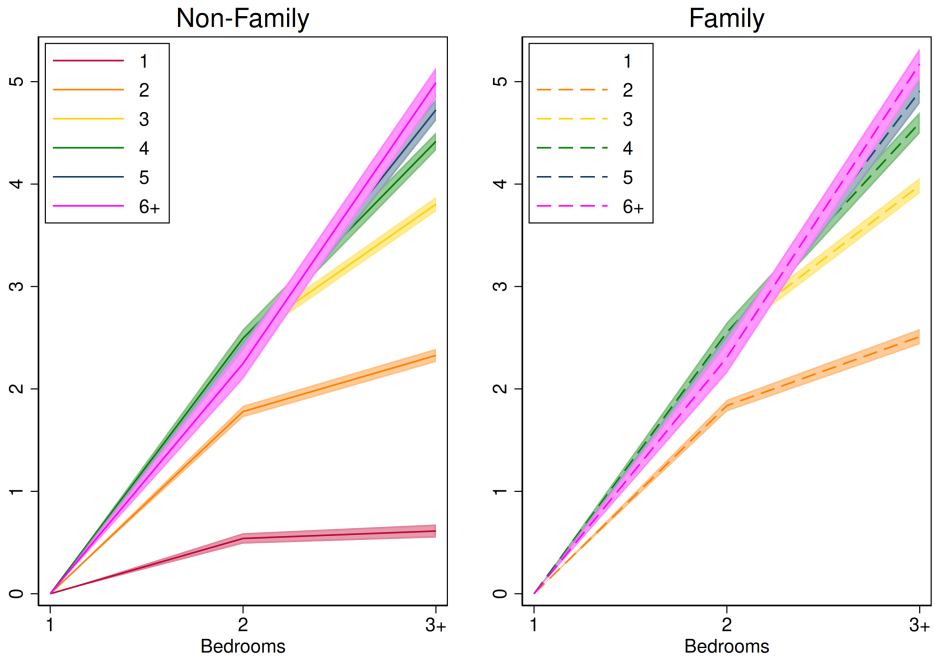


Figure 12

I assume that the value can be decomposed as:

$$v_{gh}^{zx} = u_h^{z,\tau(g,x)} + \alpha_g^{zx} + c_g^{zx}.$$

The heterogeneous housing utility that was decomposed in the previous section is $u_h^{z,\tau(g,x)}$. The living arrangement utility that will be estimated in this section is *per se*, α_g^{zx}.

The continuation value from the living arrangement choice is the discounted expectation of the value of the problem in the next period conditional on everything known in this period. This expectation is over stochastic state transition and next period values. The living arrangement choice g, along with exogenous state z and endogenous state x, determines the state transition probability distribution π. and the evolution of the endogenous state x affects the options available in the next period and hence the next period's inclusive value.

$$c_g^{zx} = \delta \mathbb{E}\left[\sum_{z'} \sum_{x'} \pi(z', x' \mid z, x, g) \times \ln\left(\sum_{h'} \sum_{g' \in \mathcal{G}(x')} \exp\left(v_{g'h'}^{z',x'} \right) \right) \middle| W_t \right]$$

Living arrangement options g include

- three options that provide $u_h^{z,\tau} = 0$ due to data constraints: group quarters, with parents, or with other family
- living in a non-family household of size s (including alone)
- having n children, and automatically converting to a family household

- converting to a homeowner household

The exogenous state z consists of time, age, and a binary immigrant variable. There are four time periods 1990, 2000, 2010, and 2020, and six ages 10-19, ..., 50-59, 60+. People aged 0-9 are dependent children born in that time period who do not make choices, and people aged 60+ either remain in that age bin in the next period or pass away.

The endogenous state x consists of

- a binary family indicator
- a binary homeowner indicator
- a binary single parent indicator
- the stock of children aged 10-19
- the stock of children aged 20-29
- the number of "extra" family members present – not the head, the spouse, or their children.

By including unmarried parents, adult children, and non-nuclear family members, the model is flexible enough to accommodate unorthodox living arrangements inherent to the Second Demographic Transition, but this introduces complexity and forces compromises in $\mathcal{G}$ and π.

- Group quarters, living with parents, and living with other family are options only available to agents that have not started families.
- Single parents are unable to have additional children. If a single parent is observed with young children, I assume that the partnership dissolution shock occurred after the children were born, and add the second parent back to the household.
- Only family households may be homeowners. A non-family homeowner household is necessarily one where the owner is the landlord and roommate to unrelated adults. Modeling this situation is beyond the scope of this paper. Luckily, there is only a small number of such households in the data, and I set them to be renters instead.
- The probability distribution over the stock of children remaining in the household in the next period is a binomial distribution with the number of trials equal to the current stock and the probability the endogenous living-with-parents probability of agents of that age. Notably, due to data constraints, this decision is treated as irreversible from the perspective of the parent, but reversible from the perspective of the child. However, because most people aged 10-19 live with their parents, the number of children of that age in the household is very often equal to the number of children of that age that the parents have ever had, and so the discrepancy for the distribution over the stock of children aged 20-29 is minimal. Older children are treated as extra (non-nuclear) family members regardless of their relationship to the head.
- Once all of a family's children have moved out, its agents automatically revert to non-family, even if the couple is married.
- Because the existence of family ties outside of cohabitation is unobserved, living with other family is treated as an option that is available to all. The distribution over the number of extra family members is a function of the number of agents who choose to live with other family, the state of the receiving household, and the distribution of households over states.

- Because there are family households that do not include children or partners in the data, non-family households may also receive extra family members and thereby become a family household just for that period.

7.2 Data

I use data from the 1990 Census, 2000 Census, 2008-12 pooled 5-year ACS, and 2018-22 pooled 5-year ACS. Although the data is cross-sectional, I am able to infer information about agents' state 10 years prior when they made their living arrangement choice. I infer that they had the same children, but 10 years younger, or that they were single if they had no children 10 or older; that homeowners were homeowners 10 years prior if they have no mortgage or if they have resided at that address for more than 10 years (and otherwise assume that homeowners were renters 10 years prior if neither of these conditions hold); and I assume that singles were also singles years prior.

Alternative living arrangements are increasingly supplanting the traditional nuclear family structure: adult children remaining in the household, divorces, and other family members or friends residing in a household. In addition to their effects on housing and living arrangement situations, it is important to accommodate these alternative structures in the model so that we are not compelled to drop or re-code increasingly large proportions of the sample. To that end "extra" family members are the most difficult to handle, but their inclusion in the model is also very important in covering as much of the sample as possible. As a result, I use 97.33% of households in the raw data, and 95.07% of agents.

Unfortunately, censuses across the world tend to under-count young children (Census Bureau, 2016)[20], making it difficult to use Census Bureau products to study fertility. To address this problem, I (1) construct age-specific 10-year birth rates for women in each age bin in each decade of study using official CDC fertility statistics and intercensal population estimates, and (2) re-weight the microdata to match both the total number of implied births for household heads of each age in each decade, as well as age-decade population totals, so that the estimation sample matches official fertility statistics and population aggregates.

7.3 Estimation

I first invert $\alpha + c$ from the aggregated choice probabilities:

$$\alpha_g^{zx} + c_g^{zx} = \ln\left(P_g^{zx} / \sum_h \exp(u_h^{z,\tau(g,x)})\right) - \ln\left(P_{g_0(x)}^{zx} / \sum_h \exp(u_h^{z,\tau(g_0(x),x)})\right),$$

where $g_0(x)$ is living alone for agents that enter the period in non-family households, and choosing to have zero children for agents entering the period as family households. It is without loss to normalize two options to have zero utility, because the value of being a family is captured in the value of having children and automatically converting to a family household.

[20] https://www2.census.gov/programs-surveys/decennial/2020/program-management/final-analysis-reports/2020-report-2010-undercount-children-mothers.pdf

Drop the h index from the value function to denote aggregation over h:

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

which is sufficient to calculate living arrangement choice probabilities and inclusive values. Assume that households forecast this object:

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

Then, to estimate c via simulation, take draws from the error distribution. For now, assume that the errors are iid and normal with a constant variance – this can later be relaxed to allow separate distributions over different indexes if the evidence points in this direction. For each draw, calculate the period-ahead inclusive value:

$$L^{zx} = \ln\left(\sum_{g' \in \mathcal{G}(x')} \exp\left(\hat{v}_{g'}^{z'x'd}\right)\right).$$

Additionally, for each draw, compute the endogenous components of π. Compute π(child stays) as a binomial distribution using the stock of children in the younger age bin and the endogenous probability that they will live with their parents after aging into the focal age bin:

$$Pr(n_{10-19} = k|\tilde{z}, x, t+1, d) = \left(\frac{n_{0-9}!}{k!(n_{0-9}-k)!}\right)(P_{\text{par}}^{(10-19,t+1,d)})^k(1 - P_{\text{par}}^{(10-19,t+1,d)})^{(n_{0-9})-k},$$

where $P_{\text{par}}^{(10-19,t+1,d)}$ is the period-ahead choice probability of given a draw d, and likewise for $Pr(n_{20-29} = k|\tilde{z}, x, t+1, d)$.

However, π(extra arrives) is a more complicated function of P_{extra}. First, combine the choice probability with the number of agents in each state, and total this up to obtain the total number of agents choosing to live as extras at time t:

$$N_{\text{extras}}^t = \sum_{\tilde{z}} \sum_x N^{\tilde{z}xt} P_{\text{extra}}^{\tilde{z}xt}.$$

Then, use a time-invariant distribution of extras over receiving states to obtain the total number of extras living with households of each "extras group"[21] $w(\tilde{z}, x)$:

$$N_{\text{extras}}^{w(\tilde{z},x),t} = s^{w(\tilde{z},x)} N_{\text{extras}}^t \quad \text{where} \quad s^{w(\tilde{z},x)} = \frac{\sum_t \sum_{(\tilde{z}',x'):w(\tilde{z}',x')=w(\tilde{z},x)} N^{\tilde{z}'x't} e(x')}{\sum_t \sum_{(\tilde{z}'',x'')} N^{\tilde{z}''x''t} e(x'')}$$

Finally, assume every agent in an extras group has an equal probability $1/N^{w(z,x)}$ of receiving

[21]The extras group is a lower-dimension function of the state, comprised of age, immigrant/native status, whether or not the agent would have been in a family household absent an extra, and whether the agent is in a household that started the period as a homeowning household.

every extra to obtain a binomial distribution:

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

To recap, I take the total number of people choosing to be an extra, apportion them among groups w according to the data over the whole sample, and then assume that each agent in a given group has an equal probability of receiving one of the extras assigned to that group. In this way the model accommodates alternative living arrangements that are themselves endogenous to the housing market even though non-resident family ties are not observed in the data.

However, this structure imposes difficulties for estimation and counterfactuals. Agents are making forward-looking decisions on the basis of the likelihood of hosting extras in the next period, which requires knowledge of – or a probability distribution over – the distribution of agents over states in the next period, which itself is a function of the transition matrix. One option would be to iterate, allowing an estimated π to feed into the next estimate. However, the distribution of agents over states depends not only on endogenous fertility but also mortality and net migration, rendering such a principled approach even more challenging.

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

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:

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

Agents then approximate

$$N^{\tilde{z}x,t+1} \approx \tilde{N}^{\tilde{z}x,t+1} = \hat{\psi}_0 N^{\tilde{z}xt} + \hat{\psi}_1^{\tilde{z}x} + \hat{\psi}_2^{\tilde{z}x}t.$$

Simulation draws for value function errors are used in the period-ahead choice probabilities:

$$\tilde{N}_{\text{extras}}^{\tilde{z}x,t+1,d} = s^{\tilde{z}x}\sum_{\tilde{z}'}\sum_{x'}\tilde{N}^{\tilde{z}'x',t+1}P_{\text{extra}}^{\tilde{z}x,t+1,d}.$$

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

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

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

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

I do not present results because they are not meaningful given the normalization on housing utility. There is no housing option that is chosen by at least one agent of every type in every year, and so there is no suitable housing outside option. This is not a problem for the estimation of housing demand parameters because of the presence of type fixed effects[22], which allow other parameters to be estimated with the log choice share on the left hand side[23] as if an option had been normalized to have zero utility. But it still means that there does not exist a normalized housing utility, and so the housing utility aggregator $\sum_h \exp(u_h^{\tau(g,x)})$ for a given τ has some unknown constant added to the utility of all options. This constant is then a part of every $\alpha_g^{zx} + c_g^{zx}$ that results from inverting P_g^{zx} using the housing utility aggregator. The α_g^{zx} that is estimated then has no inherent meaning.

However, the purpose of estimating α_g^{zx} is not to analyze the parameter estimates, but to find non-housing parameters that perfectly fit living arrangement and fertility choice probabilities, allowing non-housing utility to be held constant as housing utility changes in counterfactual exercises. The presence of a contaminating constant in α_g^{zx} is immaterial to this purpose, which we turn to next.

8 Decomposition

8.1 Setup

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, the effects of its trends are far-reaching. 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 exercise in this section sets the rent of housing units of every size and location to be equal to its value in 1990. 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 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. However, because this decomposition controls prices, it can also be thought of as a partial equilibrium exercise.

8.2 Equilibrium Definition and Solution Algorithm

The solution algorithm employed in this decomposition is also the workhorse for the counterfactuals in the next section. A change in rents r affects housing utility u and thus value functions v. This shift in v moves forecasted v', the continuation value c, and then feeds back into v. The shift in forecasted v' moves expected period-ahead choice probabilities P', transition matrix

[22] More specifically, city-type and bedroom-type fixed effects that are collinear with type fixed effects.

[23] In this case, with the un-logged choice share on the left hand side of a Poisson maximum likelihood specification, which has the same property.

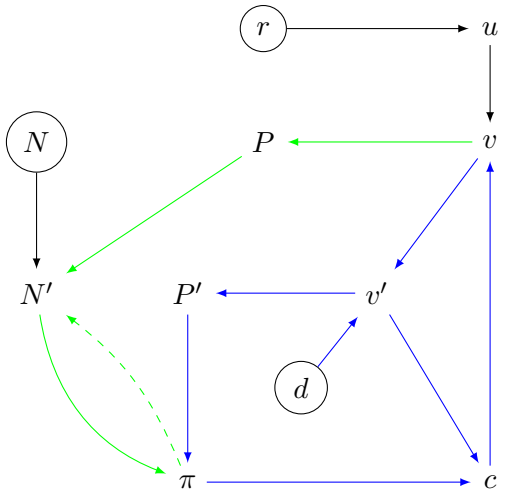


Figure 13

π (through children's decisions), the continuation value c, and feeds back into v in a second pathway. Finally, there is the complicated treatment of extra, non-nuclear family members. There is a pathway from v through v', P', π, c, and back to v that is analogous to the one for endogenous children's decisions. There is also a pathway from v, through P, to the simultaneous relationship between N' and π, to c and then back to v.

Because the current-period distribution of agents over states N affects the next-period distribution of agents over states N', and N' feeds into v, the model must be solved moving forward in time. I also solve for exogenous changes in N due to net immigration and mortality using the model baseline scenario, and include these exogenous changes in the decomposition scenario as well.

The complete set of dependencies is depicted in Figure 13. Every arrow is an equilibrium condition. Inputs r, d, and N are circled. The inner loop is shown in blue and the outer loop is shown in green.

Given rents and thus housing utilities, a partial equilibrium is then a set of value functions $v_g^{\tilde{z}xt}$ such that (where each item is a previously-stated equation):

- the definition of v_g^{zx} is satisfied
- future values $v_g^{\tilde{z}x,t+1}$ are a linear function of current values and an error term
- current choice probabilities are a logit function of current values
- future choice probabilities are a logit function of future values
- future child stay probabilities are a binomial distribution using the endogenous live-with-parents probabilities and stock of children
- future distribution of agents over states is an aggregative function of the current distribution, the transition probabilities, and current choice probabilities
- future total number of extras for each receiving type is an aggregative function of the future distribution of agents over states, the future extra choice probability, and a time-invariance constant reflecting the tendency of some household types to receive more extras

- future extra arrive probabilities are a binomial distribution using the future total number of extras for each receiving type and the total number of households of each receiving type
- the probabilities in the transition matrix are the product of endogenous child stay probabilities and extra arrive probabilities, and exogenous partnership dissolution probabilities
- the continuation value is the expectation (over future value error terms) of the product of transition probabilities and future inclusive values.

The decomposition solution algorithm is as follows. In 2000, initialize v with the new u (from the shift in r) and the estimated c. Take D draws from the error distribution. In an inner loop, update v', P' (using the draws), π, c, and v, iterating until convergence. In an outer loop, use the new v and the input N to update P, N', π, c, and v, iterating until convergence. Because the outer loop is updating π through the distribution of agents over states and not future choice probabilities, it does not require simulation. Because π has been changed, the equations of the inner loop no longer hold, so repeat the inner loop in the second outer loop iteration, and iterate all loops until convergence. Although this algorithm allows N' to affect π but not vice versa in a single outer loop iteration, this equation holds after convergence because the updated π will feed back into N' in the next outer loop iteration. After convergence, take N' and the exogenous change in the number of agents in each state to form N for 2010, and repeat. After convergence for all years, calculate the number of children born n and the total fertility rate and other outcomes.

After convergence, the number of children born and the total fertility rate are calculated as follows.

$$n_{0-9}^{t} = \sum_{\tilde{z}:t(\tilde{z})=t} \sum_{x} \sum_{g \in \mathcal{G}(x)} N^{\tilde{z}xt} P_g^{\tilde{z}xt} n(g)/2.$$

Adding over all states (in that year) and living arrangement choices, take the product of the number of agents in the state and making the choice $N^{\tilde{z}xt} P_g^{\tilde{z}xt}$ and the number of children resulting from the choice $n(g)$, and divide by two because it takes two people to make a child (and because there is no gender in the model).

$$\text{TFR}_t = \sum_{a} \frac{\sum_{\tilde{z}:a(\tilde{z})=a,t(\tilde{z})=t} \sum_{x} \sum_{g \in \mathcal{G}(x)} N^{\tilde{z}xt} P_g^{\tilde{z}xt} n(g)/2}{\sum_{\tilde{z}:a(\tilde{z})=a,t(\tilde{z})=t} \sum_{x} N^{\tilde{z}xt}}.$$

We calculate the number of children born to people in a given age, divide by the number of people of that age, and add over all ages to find the total number of children born to a hypothetical person who lives their entire life at time t. Again, notably, this object differs from standard total fertility rates which are calculated using the female population, but without gender in the model I must assume that everyone has a 50% probability of being female (and because standard TFRs use 5-year age bins, and I use 10-year age bins).

8.3 Results

The total number of children born in the data, in the decomposition, and in a reduced-form exercise are depicted in Figures 14 and 15. The reduced-form exercise uses the preferred estimate from regressions of general fertility rates on log 3+ bedroom rents, multiplies[24] it by 10 (to make

[24]This coefficient multiplied by 10 is smaller in absolute value than an identical specification where children aged 0-9 per woman aged 15-44 is regressed on log 3+ bedroom rents, likely because a family's housing demand is greater

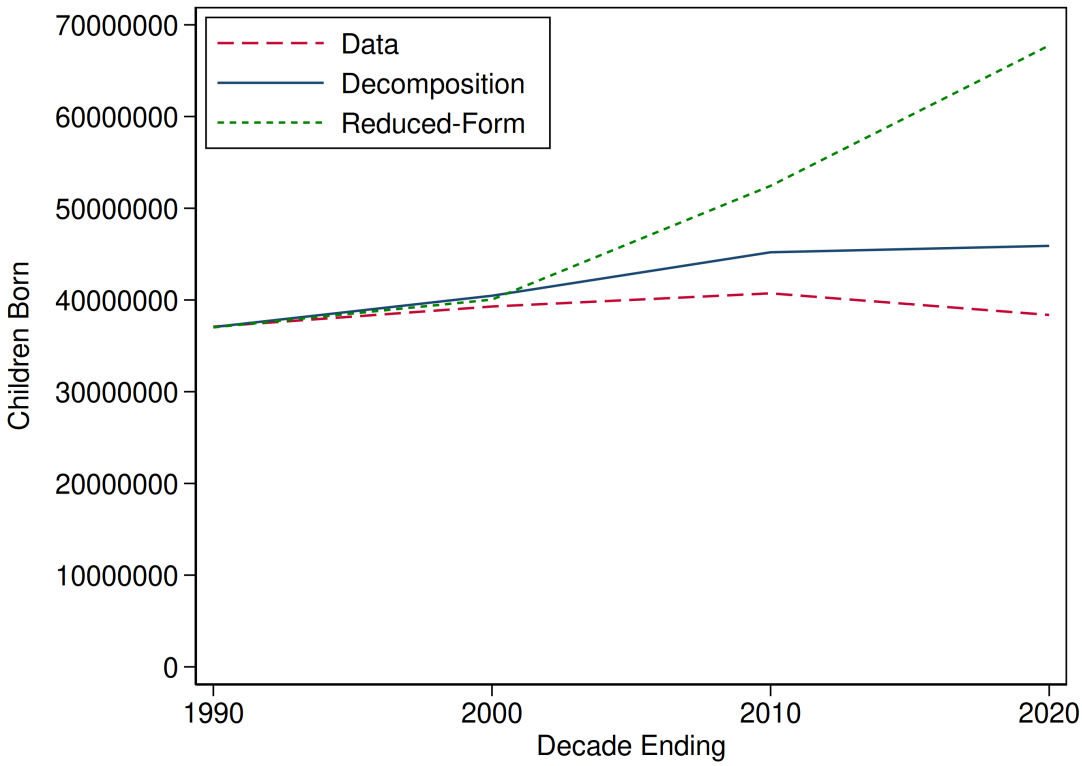


Figure 14

it comparable to the model, where rents affect the number of children aged 0-9), and applies the predicted general fertility rate to the number of women aged 15-44. The number of women aged 15-44 is allowed to evolve in the reduced-form exercise (to make it analogous to the model-based decomposition) under the assumption that new children born are 50% female, and there is a 50% probability of a woman being born in the first half or latter half of a decade.

Because the rent shock starting in 2000 is unexpected, 1990 fertility is unchanged. Because there was not much rent growth in the 1990s in the data, fertility in 2000 is similar in the decomposition and reduced-form. However, rent growth accelerated in the 2000s and 2010s, and so the gaps between the data and the decomposition and reduced-form widen. Whereas the number of children in the data declines between 2010 and 2020, in the decomposition and reduced-form it increases, albeit only slightly in the decomposition.

In the latter three decades, 13 million more children are born in the decomposition, an increase of 11% over the 118 million children born in these decades in the data. From this we can conclude that housing policy can indeed have a material effect on the number of children born and contribute substantially to ameliorating the fertility crisis. In contrast, in the reduced-form exercise a whopping 42 million more children are born – 30 million in the last decade alone – representing an increase of 36%. From this we can conclude that the sorting endogeneity

when it has 9-year-old than an infant, strengthening the sorting endogeneity that the model is intended to correct. To avoid strawmanning the reduced-form, I instead multiply the less-endogenous infant-based coefficient by 10.

decade_ending	children_data	children_counter	children_rf
1990	37 million	37 million	37 million
2000	39 million	40 million	40 million
2010	41 million	45 million	52 million
2020	38 million	46 million	68 million

Figure 15

contained in the reduced-form estimates that the model is intended to correct is substantial; although the impact of housing is material, policy counterfactuals based on such reduced-form estimates that do not account for agents' re-optimizing housing choices would severely mislead about the potential of housing policy reform in solving the fertility crisis.

Total fertility rates in the data and decomposition are shown in Figures 16 and 17.[25] Again, there is no change in 1990, and thereafter a growing difference as the gap between rents in the data and in the decomposition grow. In the decade ending 2010, TFR in the data grew even as rent growth accelerated, in part because of the growing economy, and likely fueled by homeownership and home equity wealth. The impact of housing on fertility is multidimensional, and if user costs had not also grown then the TFR would have climbed even higher to 2.16, above the replacement rate. However, the steep decline in fertility in the 2010s was mostly driven by non-housing factors, such that the decomposition TFR declines to 2.06. Overall, rising housing costs play a significant role in the U.S.'s inability to achieve replacement-rate fertility in 2020, as an increase from 1.74 in the data to 2.06 in the decomposition erases 89% of the gap between the data and replacement rate.

9 Counterfactuals

9.1 Setup

The main counterfactual exercise of this paper targets the heterogeneous intercepts of the housing supply equation via a housing development subsidy and/or regulatory tax relief. Standard YIMBY policy is to build small units and lower rents in already-dense neighborhoods, and there are at least two reasons to believe that this may be pro-natal: it would divert the demand of non-family roommate households away from larger units, and it would reduce the housing cost of fertility by diverting adult children and extra family members away from family households. On the other hand, a more directly pro-natal housing policy would be to build more and lower the rent of large units.

This counterfactual requires the addition of a supply side to the model. The supply curves are functions of the local 10-year elasticities γ_j estimated by Baum-Snow and Han (2024), which are heterogeneous over j and homogeneous otherwise. I convert these 10-year elasticities to cross-sectional elasticities by demeaning over tracts within a bedroom-year and estimating a

[25] Because TFRs are not observed at the tract level, there is no reduced-form analogue, although a decomposition-based GFR could be calculated.

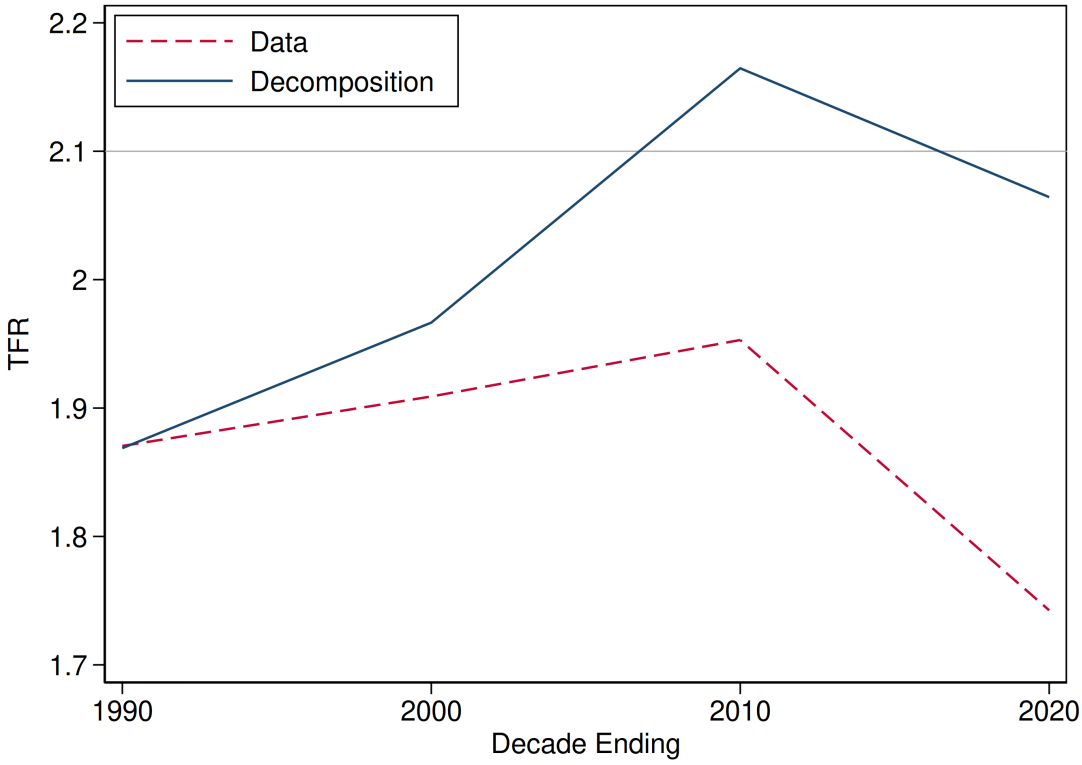


Figure 16

decade_ending	tfr_data	tfr_counter
1990	1.87	1.87
2000	1.91	1.97
2010	1.95	2.16
2020	1.74	2.06

Figure 17

shrinkage parameter $\tilde{\gamma}$ and heterogeneous intercepts with the following model:

$$\ln r_{jbt} - \ln \bar{r}_{bt} = \gamma_{jb} + \gamma_{bt} + \gamma_{jt} + \tilde{\gamma}\big(\gamma_j^{-1}(\ln H_{jbt} - \ln \bar{H}_{bt})\big) + \epsilon_{jbt}.$$

The counterfactual targets γ_{bt}, providing percentage subsidies to the construction of differently-sized units at different times:

$$\{\gamma_{bt} : b = 1, t >= 2000\} \quad \text{vs} \quad \{\gamma_{bt} : b = 3+, t >= 2000\}.$$

The quantity in this supply equation is H_{jbt}, the number of occupied housing units, which is equivalent to the number of households (as per Census Bureau definitions). This is computed as[26]

$$H_{jbt} = \sum_{\tau} H_{jbt}^{\tau}.$$

9.2 Equilibrium Definition and Solution Algorithm

Relative to the partial equilibrium definition in the decomposition sections, this equilibrium definition adds only a pathway from N and P to H, then to r, then back to v.

An equilibrium is a set of value functions v_g^{zx} and rents r_{jbt} such that:

- all conditions in the partial equilibrium definition hold
- the number of households (equivalent to the number of housing units) are an aggregative function of the number of agents as described above
- the housing supply equation holds
- housing utility is a linear function of rent
- (worth repeating) the definition of v_g^{zx} holds, where housing utility enters via the housing utility aggregator.

These additional conditions are depicted in Figure 18. Now, γ is an input and r is not. There is now an outer loop where P (and the input N) determine H, which feeds back into r via the housing supply equation, and u and v.

9.3 Results

10 Conclusion

Fertility is an individual decision, which, when aggregated over all individuals in a society, has societal consequences. Declining fertility in developed economies portends a demographic crisis with variegated consequences that cannot be indefinitely forestalled with immigration. Because fertility is an individual decision, it responds to incentives. Since children require additional housing, it follows that rising housing costs have contributed to declining fertility, and may be an effective lever to arrest or reverse the decline.

[26]

$$H_{jbt}^{\tau(g,x)} = \begin{cases} N_{jbt}^{\tau(g,x)}/\text{size}(x) & \text{if} \quad \text{family}(x) = 0 \;\&\; g \notin \{\text{GQ, Par, Extra}\} \\ 0 & \text{if} \quad g \in \{\text{GQ, Par, Extra}\} \\ N_{jbt}^{\tau(g,x)}/2 & \text{if} \quad \text{family}(x) = 1, \end{cases}$$

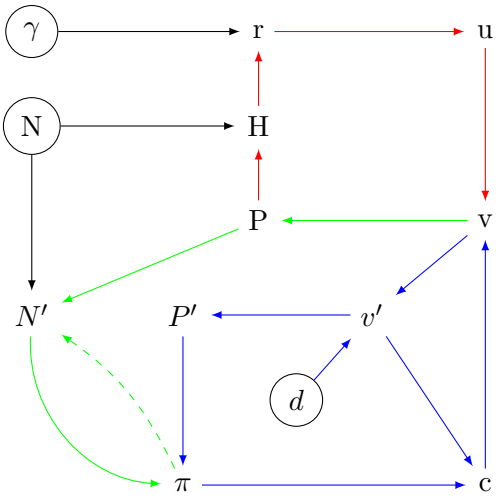


Figure 18

A large literature has examined how income affects fertility, and another literature has shown that housing affects fertility in an intuitive way as a major budget component and source of wealth. However, housing is an endogenous choice that is made jointly with fertility with dynamic life-cycle considerations in mind. Almost all papers about the relationship between housing costs and fertility use geographic variation, which cannot fully adjust for sorting endogeneity regardless of how strong the causal identification strategy is otherwise due to a SUTVA violation.

To build on the existing reduced-form literature, this paper specifies a dynamic model of joint fertility and housing choice to examine how the rising costs of both large and small housing units affect fertility decisions. Agents make a dynamic living arrangement decision – which along with fertility includes being single, living with roommates, living with parents, living with other family members, and converting to a home-owning household – knowing how their static 10-year demand for housing units of various sizes in different neighborhoods will change, with rational expectations over how their choice impacts the likelihood that adult children or other family members will live with the household and shape housing demand in the future. The values of living arrangement decisions explained by everything other than housing are estimated so as to be held constant in decomposition and counterfactual exercises. Because having children increases price sensitivity and tastes for larger and more expensive units (a pattern predicted by a Beckerian quantity-quality model with housing), increases in housing costs systematically reduce the value of fertility.

To estimate the housing demand part of the model, I provide and exploit new econometric theory results about sufficient data for logit demand with observed heterogeneity using the iterative proportional fitting (IPF) algorithm, even though the joint distribution of neighborhood, unit size, household head age, tenure, family/non-family status, and household size is not observed outside of restricted-access research data centers. These results extend to some cases where IPF-generated synthetic joint distributions are used to construct left-hand side variables for OLS and 2SLS models. I also provide a new falsification test for commonly-used "donut" instrumental variables. Housing demand parameter estimate signs almost completely match theoretical predictions, patterns of heterogeneity match those predicted by theory and necessary

for the hypothesized mechanism for rising housing costs to decrease fertility, and numeraire coefficient estimates for owners closely match Bayer et al (2016), who used a completely different strategy.

In a decomposition exercise, I prevent rents from growing after 1990 – which doubles as an imposition of perfectly elastic housing supply – to form an upper-bound estimate of the impact of housing costs on fertility, finding that 11% more children would have been born and that 90% of the gap between the replacement and 2010s total fertility rate would have been erased, an effect which is much smaller than a naive application of reduced-form estimates. From this we can conclude both that the sorting endogeneity that motivates the model is a serious concern, and that even a modest housing policy reform can materially contribute to ameliorating the fertility crisis. A more realistic counterfactual exercise is in progress, which would compare the effects of standard YIMBY policy – subsidizing small units – versus a more explicitly pro-natal policy of subsidizing large units.

Funding Disclosure

This research was funded in part by a CGS-D Doctoral Award from the Social Sciences and Humanities Research Council of Canada (SSHRC) and an Ontario Graduate Scholarship (OGS).

References

[1] Gunnar Andersson. Childbearing after migration: Fertility patterns of foreign-born women in sweden. *Demographic Research*, 10(4):45–86, 2004.

[2] Kadir Atalay, Dian Li, and Stephanie Whelan. Housing wealth, fertility intentions and fertility: Evidence from australia. *Journal of Housing Economics*, 54:101744, 2021.

[3] David H. Autor, David Dorn, and Gordon H. Hanson. When work disappears: Manufacturing decline and the falling marriage-market value of young men. *American Economic Review: Insights*, 1(2):161–178, 2019.

[4] Nathaniel Baum-Snow and Lu Han. The microgeography of housing supply. *Journal of Political Economy*, 132(S1):S68–S120, 2024.

[5] Patrick Bayer, Fernando Ferreira, and Robert McMillan. A unified framework for measuring preferences for schools and neighborhoods. *Journal of Political Economy*, 115(4):588–638, 2007.

[6] Patrick Bayer, Robert McMillan, Alvin Murphy, and Christopher Timmins. A dynamic model of demand for houses and neighborhoods. *Econometrica*, 84(3):893–942, 2016.

[7] Patrick Bayer and Christopher Timmins. Estimating equilibrium models of sorting across locations. *Economic Journal*, 117(518):353–374, 2007.

[8] Gary S. Becker. An economic analysis of fertility. In Universities–National Bureau Committee for Economic Research, editor, *Demographic and Economic Change in Developed Countries*, pages 209–240. Princeton University Press (for NBER), Princeton, 1960.

[9] Gary S. Becker and H. Gregg Lewis. On the interaction between the quantity and quality of children. *Journal of Political Economy*, 81(2, Part 2):S279–S288, 1973.

[10] Johanna Bergsvik, Rannveig K. Hart, Siri Vangen, and R. K. Hart. Can policy stall the fertility fall? a review of evidence. *Population and Development Review*, 47(4):913–964, 2021.

[11] Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890, 1995.

[12] Steven Berry, James Levinsohn, and Ariel Pakes. Differentiated products demand systems from a combination of micro and macro data: The new car market. *Journal of Political Economy*, 112(1):68–105, 2004.

[13] Sonia Bhalotra and Damian Clarke. The twin instrument: Fertility and human capital investment. *Journal of the European Economic Association*, 18(6):3090–3139, 2020.

[14] Yvonne M. M. Bishop. *Inference and Modeling in Multidimensional Contingency Tables*. Ph.d. dissertation, Harvard University, 1967.

[15] Yvonne M. M. Bishop, Stephen E. Fienberg, and Paul W. Holland. *Discrete Multivariate Analysis: Theory and Practice*. MIT Press, Cambridge, MA, 1975.

[16] Dan A. Black, Natalia Kolesnikova, Seth G. Sanders, and Lowell J. Taylor. Are children "normal"? *Review of Economics and Statistics*, 95(1):21–33, 2013.

[17] Francine D. Blau. Immigration and female labor supply in the united states. NBER Working Paper 21756, National Bureau of Economic Research, 2015.

[18] David E. Bloom, Michael Kuhn, and Klaus Prettner. The economic consequences of declining fertility. *Annual Review of Economics*, 16:—, 2024.

[19] Mary O'Malley Borg. The income–fertility relationship: Effect of the net price of a child. *Demography*, 26(2):301–310, 1989.

[20] David Cesarini, Erik Lindqvist, Robert Östling, et al. Wealth and fertility: Evidence from swedish lotteries. Technical report, Working paper, 2023. Lottery-based wealth shocks; Swedish administrative data.

[21] Warren Clark and Ana Ferrer. The effects of housing prices on childbearing in canada. *Economics (E-Journal)*, 13:1–33, 2019.

[22] William A. V. Clark. Do women delay childbearing in response to housing market constraints? *Demographic Research*, 27:163–198, 2012.

[23] Christopher Conlon and Jeff Gortmaker. Incorporating micro data into differentiated products demand estimation with pyblp. NBER Working Paper 31605, National Bureau of Economic Research, 2023. Revised 2024.

[24] Emma Cumming and Lisa J. Dettling. Interest rate reductions, mortgage payments, and fertility. *Review of Economic Studies*, 91(—):—, 2024.

[25] N. Meltem Daysal, David Gunnarsson, Thit Aagaard Jensen, and Lars Skipper. Home prices, fertility, and early-life health outcomes. *Journal of Public Economics*, 197:104380, 2021.

[26] W. Edwards Deming and Frederick F. Stephan. On a least squares adjustment of a sampled frequency table when the expected marginal totals are known. *Annals of Mathematical Statistics*, 11(4):427–444, 1940.

[27] Lisa J. Dettling and Melissa S. Kearney. House prices and birth rates: The impact of the real estate market on the decision to have a baby. *Journal of Public Economics*, 110:82–100, 2014.

[28] Lisa J. Dettling and Melissa S. Kearney. Mortgage innovation and the u.s. baby boom. Technical report, Working paper, 2025. WP version per author request; forthcoming journal status unknown.

[29] Matthias Doepke, Anne Hannusch, Fabian Kindermann, and Michèle Tertilt. The economics of fertility: A new era. In Shelly Lundberg and Alessandra Voena, editors, *Handbook of the Economics of the Family, Vol. 1 (2nd ed.)*, pages 151–254. Elsevier, 2023.

[30] Richard A. Easterlin. What will 1984 be like? socioeconomic implications of recent twists in age structure. *Demography*, 15(4):397–432, 1978.

[31] Belton M. Fleisher and George F. Rhodes. Fertility, women's wage rates, and labor supply. *American Economic Review*, 69(1):14–24, 1979.

[32] Amit Gandhi and Jean-François Houde. Measuring substitution patterns in differentiated-products markets. *American Economic Review*, 109(10):3400–3444, 2019.

[33] Paulo Guimarães, Octávio Figueiredo, and Douglas Woodward. A tractable approach to the firm location decision problem. *Review of Economics and Statistics*, 85(1):201–204, 2003.

[34] James J. Heckman and Rasmus Landersø. Lessons from denmark about inequality and social mobility. *Labour Economics*, 79:101945, 2022.

[35] James J. Heckman and James R. Walker. The relationship between wages and income and the timing and spacing of births: Evidence from swedish longitudinal data. *Econometrica*, 58(6):1411–1441, 1990.

[36] Melissa S. Kearney and Riley Wilson. Male earnings, marriageable men, and nonmarital fertility: Evidence from the fracking boom. *Review of Economics and Statistics*, 100(4):678–690, 2018.

[37] Aniket D. Kulkarni, Denise J. Jamieson, Howard W. Jones, Dmitry M. Kissin, Maria F. Gallo, Maurizio Macaluso, and Eli Y. Adashi. Fertility treatments and multiple births in the united states. *New England Journal of Medicine*, 369(23):2218–2225, 2013.

[38] Mark Lino, Kevin Kuczynski, Nestor Rodriguez, and TusaRebecca Schap. Expenditures on children by families, 2015. Miscellaneous Publication 1528-2015, U.S. Department of Agriculture, Center for Nutrition Policy and Promotion, 2017. January 2017; Revised March 2017.

[39] Ziqian Liu and Yu Zhang. The effect of house prices on fertility: Evidence from house purchase restrictions. Working paper, September 30, 2024 draft (AEA 2025 Program), 2024.

[40] Michael F. Lovenheim and Kevin J. Mumford. Do family wealth shocks affect fertility? evidence from U.S. home price changes. *Journal of Human Resources*, 48(1):1–35, 2013.

[41] Kevin Milligan. Subsidizing the stork: New evidence on tax incentives and fertility. *Review of Economics and Statistics*, 87(3):539–555, 2005.

[42] John Mullahy. Instrumental-variables estimation of count data models: Applications to models of cigarette smoking. *Review of Economics and Statistics*, 79(4):596–606, 1997.

[43] Yan Pan and Tian Yang. Housing prices and fertility in china: Evidence from homeowners. *International Regional Science Review*, 45(5):534–554, 2022.

[44] Daniel Parent and Ling Wang. Tax incentives and fertility: Evidence from quebec. *Canadian Journal of Economics*, 40(3):593–613, 2007.

[45] Emilio A. Parrado and S. Philip Morgan. Intergenerational fertility among hispanic women: New evidence of immigrant assimilation. *Demography*, 45(3):651–671, 2008.

[46] Amil Petrin. Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy*, 110(4):705–729, 2002.

[47] Curtis J. Simon and Robert Tamura. Do higher house prices reduce fertility? *Regional Science and Urban Economics*, 39(4):483–513, 2009.

[48] Yung-Yu Tsai, Hsing-Wen Han, Kuang-Ta Lo, and Tzu-Ting Yang. The effect of financial resources on fertility: Evidence from administrative data on lottery winners. *arXiv*, 2023.

[49] United Nations, Department of Economic and Social Affairs, Population Division. World population prospects 2022: Summary of results. Technical Report ST/ESA/SER.A/448, United Nations, New York, 2022.

[50] U.S. Census Bureau. 2020 census research and testing: Investigating the 2010 undercount of young children—examining the coverage of young mothers. Internal research memorandum, U.S. Census Bureau, 2016. 2016.08.

[51] Bernardus van Doornik, Dimas Fazio, Tarun Ramadorai, and Jānis Skrastiņš. Housing and fertility. Working Paper 612, Banco Central do Brasil, December 2024.

[52] Evgeny Yakovlev and Sergey Sorvachev. Fertility effects of russia's maternity capital program. Technical report, Working paper, 2025. Working paper on Russia's "Maternal (Maternity) Capital" policy.

[53] Nishant Yonzan, John Bluedorn, Luca Antonio Ricci, Malhar Nabar, et al. Universal cash transfers and fertility: Evidence from alaska's permanent fund dividend. Imf working paper, International Monetary Fund, 2024.

A Microeconomic Theory: 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 $z = (q, c, h),\ z' = (q', c', h')$

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

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

$$U(n, s) = \mathcal{U}(n, z_s)$$

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

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

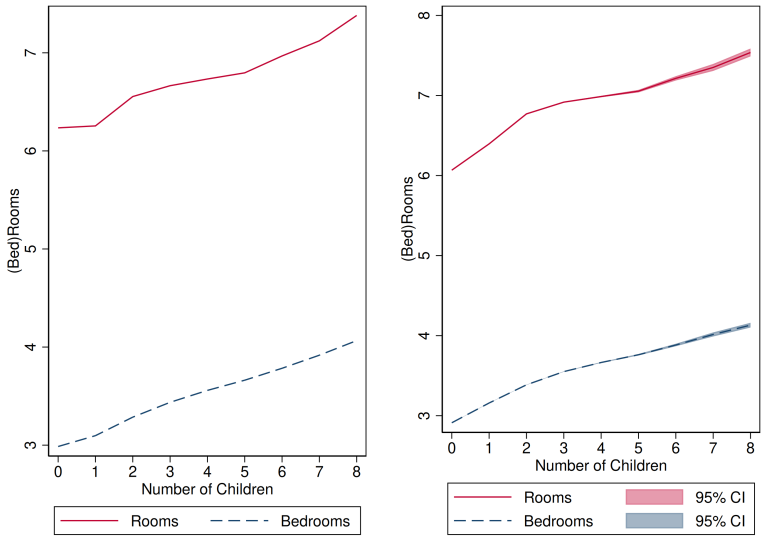
and hence we may maximize $\mathcal{U}$ by choosing n to maximize U directly and through its impact on the inner problem's value.

B Empirical Evidence: Larger Families Consume More Housing

The core mechanism of this paper is that larger families have stronger preferences for living space and rent. The applied theory model showed that the former is a driver of the latter, as fertility causes an increase in housing consumption that results in less income to spend on child quality and consumption. To observe in the data that fertility causes an increase in housing consumption validates the idea that larger families are more rent-sensitive.

To this end, I present aggregates and non-parametric OLS results, and then I use plausibly exogenous variation in fertility induced by multiple births (twin, triplet, *etc.*). Again, I use Census and ACS data. I measure housing consumption with the number of rooms and bedrooms in the house occupied by the family, and I measure fertility with the number of children in the household. I restrict the sample to family households where a spouse is present to make more comparable households with zero and a positive number children. I use city-year-age fixed effects and heteroskedasticity-robust standard errors throughout.

Aggregate (left) and non-parametric OLS (right) results evaluated at the average city-year-age fixed effect are presented below. The OLS results closely match the aggregates, with a positive monotonic relationship between house size and the number of children.



It is possible that living in a larger dwelling may induce fertility in a form of reverse causality, and so I next use a form of plausibly exogenous variation in fertility in multiple births. Multiple births historically have occurred due to uncontrollable events in the female reproductive system and thus could not be chosen, making multiple births a popular source of exogenous variation in fertility. However, the rise of assisted reproductive technology has moved multiple births from a random event to a choice variable for families with the necessary resources, such that about 23% of twin births occurred with some form of medical assistance in 1997 (Kulkarni et al, 2013). Additionally, it has lately been suggested that maternal health is correlated with successfully bringing a multiple pregnancy to term via selective miscarriage (Bhalotra and Clarke, 2020). To mitigate these concerns, for these regressions I restrict the sample to 1990, the earliest time period that I study, and I include log household income as a control variable.

I calculate "extra children" as the number of children in the household that have the same age, minus one for however many duplicated ages there are. Then, "birth events" are the number of times that a member of the household gave birth, and together with "extra children" sum to the total number of children in the household. Birth events are taken to be an endogenous choice, but extra children are taken to be exogenous. I regress the number of rooms and the number of bedrooms on both of these measures, in addition to separate specifications using the total number of children.

The results of these regressions are presented in the table below. In both cases, extra children has a smaller coefficient than total children, but the coefficient is still positive and very statistically significant. In addition to the theoretical and structural evidence presented elsewhere in the paper, there is also reduced-form evidence for the idea that larger households are more rent-sensitive.

	(1) rooms	(2) rooms	(3) bedrooms	(4) bedrooms
Children	0.247***		0.185***	
	(0.00118)		(0.000647)	
Extra Children		0.180***		0.153***
		(0.00984)		(0.00535)
Birth Events		0.249***		0.186***
		(0.00120)		(0.000658)
Ln HH Inc.	0.889***	0.889***	0.399***	0.399***
	(0.00232)	(0.00232)	(0.00117)	(0.00117)
N	2172211	2172211	2172211	2172211

Standard errors in parentheses

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

C Econometric Theory: Extended Proofs

C.1 Proof of Theorem 1

Let $N_j^{as} = \exp(\lambda^{as} + \lambda_j^a + \lambda_j^s + \tau_j^{as})$, where each λ may be a fixed effect or a coefficient on an explanatory variable. For dependent variable Y_j^{as}, the log-likelihood is

$$\mathcal{L} = \sum_j \sum_a \sum_s Y_j^{as}(\hat{\lambda}^{as} + \hat{\lambda}_j^a + \hat{\lambda}_j^s) - \exp(\hat{\lambda}^{as} + \hat{\lambda}_j^a + \hat{\lambda}_j^s).$$

Assuming $\hat{\lambda}_{j'}^{a'}$ is a fixed effect, differentiate the log-likelihood with respect to it:

$$\frac{\partial \mathcal{L}}{\partial \hat{\lambda}_{j'}^{a'}} = \sum_s Y_{j'}^{a's} - \exp(\hat{\lambda}^{a's} + \hat{\lambda}_{j'}^{a'} + \hat{\lambda}_{j'}^s).$$

Now compare the gradient for $Y_j^{as} = N_j^{as}$ vs $\hat{N}_j^{as}$:

$$\frac{\partial \mathcal{L}^N}{\partial \hat{\lambda}_{j'}^{a'}} - \frac{\partial \mathcal{L}^{\hat{N}}}{\partial \hat{\lambda}_{j'}^{a'}} = \left(\sum_s N_{j'}^{a's}\right) - \left(\sum_s \hat{N}_{j'}^{a's}\right) = N_{j'}^{a'} - N_{j'}^{a'} = 0.$$

Assuming $\hat{\lambda}_{j'}^{a'}$ is a coefficient on an explanatory variable $\hat{\beta}^{a'} x_{j'}$, differentiate the log-likelihood with respect to $\hat{\beta}^{a'}$

$$\frac{\partial \mathcal{L}}{\partial \hat{\beta}^{a'}} = \sum_j \sum_s \left(Y_j^{a's} - \exp(\hat{\lambda}^{a's} + \hat{\beta}^{a'} x_j + \hat{\lambda}_j^s)\right) x_j.$$

Now compare the gradient for $Y_j^{as} = N_j^{as}$ vs $\hat{N}_j^{as}$:

$$\frac{\partial \mathcal{L}^N}{\partial \hat{\beta}^{a'}} - \frac{\partial \mathcal{L}^{\hat{N}}}{\partial \hat{\beta}^{a'}} = \sum_j \sum_s x_j (N_j^{a's} - \hat{N}_j^{a's}) = \sum_j x_j (N_j^a - N_j^a) = 0.$$

Identical arguments apply to the other λ, and so for models of this family the gradients and maximizers are identical regardless of whether N_j^{as} or $\hat{N}_j^{as}$ is used because they are functions of the marginal distributions alone.

C.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:

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

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:

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

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

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

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:

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

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.

C.3 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

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

and so the log-likelihood and gradient are

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

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

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

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

$$Z'\hat{\omega} = \vec{0},$$

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.