Residential Segregation and the Demand for Schooling

Motaz Al-Chanati[†]

November 2021

JOB MARKET PAPER

COMPRESSED VERSION

Click here for latest (and full-size) version

Abstract

How do schools affect where families choose to live and does their effect contribute to residential segregation? I study these questions using unique administrative microdata from Auckland, New Zealand, an ethnically diverse – but segregated – city. I develop and estimate a dynamic model of residential choice where forward-looking families choose neighborhoods based on their children's schools, local amenities, and moving costs. Previous studies typically estimate school quality valuations using a boundary discontinuity design. I leverage attendance zones in this setting to also generate reduced form estimates using this methodology. The structural model estimates show that the valuation of school quality varies by the child's school level and the family's ethnicity; the reduced form approach, however, cannot capture this heterogeneity. Moreover, I find that the reduced form estimates are aligned only with white families' valuations of quality. The model estimates also show that families experience a high disutility from moving houses if it results in their child changing school. In counterfactuals, I show that residential segregation increases as the link between housing and schools weakens.

Keywords: *residential choice, school quality, dynamic model, neighborhood sorting, housing markets*
JEL Classification: *C55, I20, J10, R21, R31*

[†]Columbia University, Department of Economics. Email: mra2165@columbia.edu.

Acknowledgements: I am grateful to Bernard Salanié, Miguel Urquiola, Cailin Slattery, Sandy Black, Vinayak Iyer, and Maggie Shi for their valuable advice and comments. I am also grateful to the participants at Columbia's Applied Research Methods, Applied Micro Theory Colloquium, IO Colloquium, and Trade and Urban Colloquium for their feedback. Finally, I am especially thankful to the staff at Stats NZ for all their work in providing access and support for the data lab.

Disclaimer: These results are not official statistics. They have been created for research purposes from the Integrated Data Infrastructure (IDI) which is carefully managed by Stats NZ. For more information about the IDI please visit https://www.stats.govt.nz/integrated-data/.

1 Introduction

Residential segregation is a prevalent issue in cities across the world (Logan and Stults, 2021; Boterman et al., 2021; Jones et al., 2018). It often reflects barriers to mobility (Bergman et al., 2019), harms to long-term well-being (Kramer, 2018), and racial disparities (Logan and Parman, 2017). Understanding what drives residential *choice* is critical to designing policies that address residential *segregation*. Schools are a key factor for families' choice of neighborhood; the demand for schools plays out through the housing market as public schools typically base admissions on the student's home address (OECD, 2019). There is also consistent evidence that households are willing to pay higher housing costs to access better quality schools (Black and Machin, 2011). Moreover, schools are a particularly important amenity to study because there is a large scope for policy interventions that change the relationship between schools and neighborhoods – for example, voucher programs or expanding school choice.

In this paper, I study how schools affect where families choose to live. By understanding this, I can then evaluate whether schools contribute to residential segregation. To study this link, I develop a dynamic model of within-city neighborhood choice that incorporates children's schooling into the family's payoffs. Each period, families decide whether to stay in their current location or move to a new neighborhood. The value of a neighborhood is affected by the quality of its schools and the ethnic composition of its residents. Families are also forward-looking: they forecast that neighborhood amenities change over time, that new children may be born, and that their preferences for schools change as their children enter and exit each school level (elementary, middle, and high school). I also account for costs to moving, both in terms of changing one's physical home as well as a cost to changing a child's school. The forward-looking behavior and the moving costs give rise to the dynamic forces in the model.

Residential choice is best analyzed through a dynamic framework, as reflected in a burgeoning empirical literature on estimating dynamic residential choice models (Bayer et al., 2016; Diamond et al., 2017; Davis et al., 2018; Caetano, 2019; Caetano and Maheshri, 2019; Almagro and Domínguez-Iino, 2021). However, with the exception of Caetano (2019), these studies have largely ignored the role of children.[1] I present descriptive evidence that children affect the family's moving decisions and that these patterns are consistent with dynamic behavior. These forces are therefore not being captured in the existing empirical literature on residential choice.[2] A key reason for the omission of children is that the data required to estimate such a model is often difficult to obtain. Studying residential segregation also requires data on a heterogeneous sample; in contrast, much of the existing evidence is estimated using

[1] I use the identification strategy proposed by Caetano (2019) in this paper. Our papers are distinguished in the way we model children and our resulting estimation strategies. Additionally, I aim to study residential segregation, while his paper studies the returns to financing public schools.

[2] There is a theoretical literature on schools and neighborhood choice, such as Epple and Romano (1998), Nechyba (2003) and Avery and Pathak (2021). Zheng (2021) empirically shows that parental valuation of quality changes as charter schools enter and uses this to give suggestive evidence of subsequent neighborhood sorting. This past literature motivates the importance of the children/school channel. The cost to changing schools is, however, an understudied topic.

data with a disproportionately high share of white households.[3]

This paper overcomes these challenges by studying Auckland, New Zealand. This city of 1.5 million is ethnically diverse: its population is 42% white, 31% Asian, and 25% Polynesian. It is also highly segregated; by some measures, it is as residentially segregated as Los Angeles and New York. Through the New Zealand statistics agency, I have access to a variety of administrative datasets including national censuses, education records, and birth records. I use this microdata to create a unique panel dataset of families that tracks their home locations and children's schools over a 12 year period. I provide causal evidence that motivates incorporating schools into a residential choice model: school quality causes house prices to increase, and a child's grade has a causal effect on the likelihood of a family moving. Estimating this model using standard techniques (e.g. Hotz and Miller (1993)) is difficult due to the high dimensionality of the problem. In my setup, I have over 6 million possible states and 44 actions that families can choose from. To overcome this, I employ the novel computationally light ECCP estimator developed by Kalouptsidi et al. (2020).[4]

With the ECCP estimator, I derive a linear equation that compares moving probabilities to the *flow* (per period) utility of living in a given neighborhood. However, there is a key identification challenge: components of the observable utility (e.g. school quality) are likely correlated with unobserved neighborhood characteristics (e.g. parks, safety). Parameter estimates by OLS will therefore be biased. I address this endogeneity problem in two ways. Using the identification strategy of Caetano (2019), I compare the moving probabilities between families with school-aged children and families without school-aged children to isolate the flow utility generated by schools. I then use an instrumental variables strategy to account for the remaining endogenous variables: house prices and ethnic shares. I use a Hausman-style instrument of non-neighboring prices and a past settlement instrument to generate simulated ethnic shares.

My analysis generates three findings. First, the valuation of school quality varies by the child's school and the family's ethnicity. I find that non-white (Polynesian and Asian) families have relatively stable preferences for quality across all school levels, with slightly lower valuations for high school. This is in contrast to white families, who value high school quality the most. I benchmark these results against estimates from the standard approach in the reduced form literature: a boundary discontinuity design (BDD). I find that the BDD estimates most closely mirror those of white families. This is likely explained by the fact that the reduced form analysis uses house sales, and that most home owners are white. This raises concerns on the generalizability of current estimates in the literature.

[3] As an example of some papers and their data: Bayer et al. (2016) in San Francisco, subset to 100% white households (53% of full dataset); Caetano and Maheshri (2019) using similar data in the Bay Area have a 72% white sample; Caetano (2019) in Minnesota, 92% white; Almagro and Domínguez-Iino (2021) in Amsterdam, 69% Dutch or Western-origin. Diamond et al. (2017) and Davis et al. (2018) both use broadly sampled data in Californian cities, but do not observe race or ethnicity.

[4] Almagro and Domínguez-Iino (2021) also use the ECCP estimator to estimate a similar model (without children) using Dutch data.

Second, I find that parents experience a high disutility to interrupting their children's education, which I term as *school moving costs*. This is captured by the child changing the school they attend as a result of the family moving.[5] Moreover, this effect is stronger the younger the child. For families of all ethnicities, the school moving cost is 1.5 times larger for elementary school than it is for high school. These costs can also be significant: for a white family to offset the cost of changing a first grader's school, the new school would need to have a quality score that is 2 standard deviations higher than the old school. Previous studies have found large moving costs for households (Bayer et al., 2016; Diamond et al., 2017), but this points to a specific source of that cost and with important policy implications. These school moving costs could hamper the effectiveness of school choice or voucher programs if families are unwilling to interrupt their child's education – even if it means attending a higher quality school.

Third, I use the model estimates to conduct simplified counterfactuals. For computational reasons, I simulate the model as a static problem; however, by calibrating the model, the baseline simulation matches observed segregation patterns in the data. I run two counterfactuals. In the first, I remove school quality from the family's utility. This represents equal school utility across neighborhoods. I find that segregation increases relative to the baseline simulation. In the second counterfactual, I remove all school-related components from the family's utility. In particular, families no longer experience a school moving cost. Here, segregation increases again, by almost double the rate of the first scenario. This change is driven primarily by an increase in Polynesian-white segregation. I find that in each successive counterfactual, white families are more likely to move. Moreover, they are more likely to move to areas with a higher share of other white families. This result is consistent with Caetano and Maheshri (2019), who find that the presence of moving costs reduces segregation.

The remainder of the paper is organized as follows. Section 2 provides an overview of the New Zealand setting and a description of the data. Section 3 gives causal and descriptive evidence, including the boundary discontinuity design. Section 4 outlines the model and the derivation of the main estimating equation. Section 5 describes the data preparation and estimation process. Section 6 provides a discussion of the results and presents the counterfactual simulations.

[5]This is cost is only experienced for school changes within the same level, but not between levels (e.g. changing from elementary to middle school).

2 Setting and Data

2.1 New Zealand Educational System

In New Zealand, there are three levels of schooling: primary (grades 1-6), intermediate (grades 7-8), and high school (grades 9-13). I focus on public schools, which are free to attend for New Zealand citizens and permanent residents. Approximately 84% of students in Auckland attend public schools.

There are three features that distinguish the New Zealand public educational system. First, it is a decentralized system where schools are independent autonomous entities (akin to charter schools in the United States). Second, it is a system of complete choice: students can apply to any public school regardless of where they live. However, some schools – but not all – have a catchment area (known as a "school zone"). Any student who lives inside a school's zone receives guaranteed admission to that school. Third, the measure of school quality is an index from 1 to 10 called the decile score. Decile scores are not based on test scores, but rather are a measure of the student body's socioeconomic status. New Zealand does not have standardized testing; consequently, decile scores are commonly (mis)interpreted as a measure of school quality. In Appendix A.1, I provide further details about these features and why they make New Zealand an ideal setting to study the research questions.

2.2 Residential Segregation in Auckland

Auckland is an ethnically diverse city.[6] 42% of its 1.5 million residents were born overseas, which is a higher proportion than in New York City (36.2%) and Los Angeles (36.4%).[7] I will focus on ethnicity as the form of heterogeneity in the population. In my analysis, I designate three distinct ethnic groups: white, Polynesian, and Asian. In Appendix A.2, I describe how I create these groupings.

2.2.1 Statistics about Ethnic Groups

The ethnic groups exhibit differences in their incomes, family structures, and housing patterns. Using microdata from the 2018 New Zealand census, I provide summary statistics on Auckland families by their ethnic group. Table A.4 shows statistics about the parents (i.e. at the individual level) and Table A.5 shows statistics at the family level. These indicate that, on average, white families are the richest group while Polynesian families are the poorest. 67% of white families reported a household income above NZ$100,000 (approximately US$73,000), while only 32% of Polynesian and 42% of Asian families reported the same. The income measure tends to correlate with other socio-economic indicators, such as whether the parents are employed or have a university degree. Asian families tend to have lower home

[6] In this paper, "Auckland" refers to the urban area of the wider Auckland region. Figure A.13 indicates this on a map.

[7] Foreign-born statistics are author's calculations from the 2018 New Zealand census (for Auckland) and the 2019 ACS 1-year estimate (for the U.S. cities)

tenure rates (mean of 6 years in current home versus 9 for the other two groups), in part due to a higher rates of immigration (18% of Asian parents immigrated in the last 5 years). Families of all ethnicities have approximately 2 children, on average.

Figure 1: School Quality versus Ethnicity

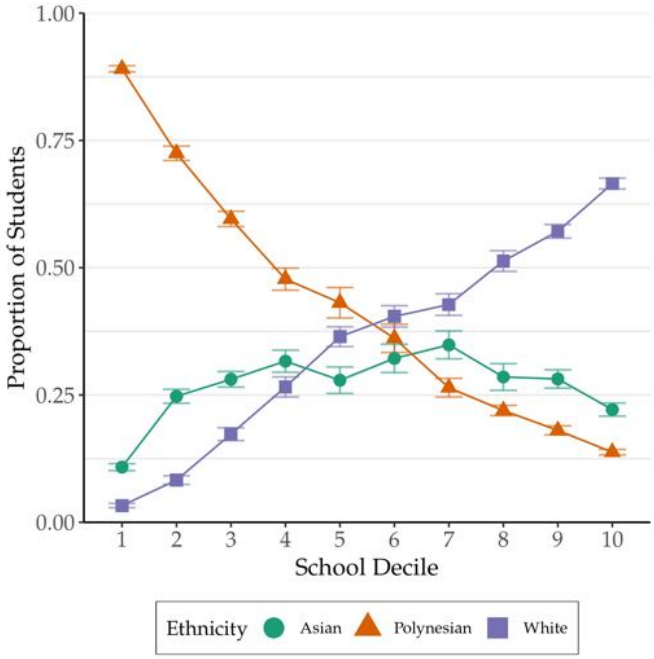

Note: The figure shows how the ethnic share of students varies with decile score (the measure of school quality) in Auckland schools for 2010-2019. I calculate the ethnic share for each school-year separately and then average values within each decile score (discrete values of 1–10). The three ethnicities shown (white, Polynesian, Asian) are distinct groupings; other ethnicities are not plotted. 95% confidence intervals are plotted with each marker.

Ethnic differences are also reflected in the children's schools. Figure 1 shows the relationship between school decile scores and the student body ethnic share. In higher decile schools, there is a higher proportion of white students and a lower proportion of Polynesian students. Similarly, there is a strong correlation between house prices and ethnic shares (Figure A.15). Areas with a higher white share (and lower Polynesian share) tend to have higher average house prices. My analysis focuses on ethnicity and residential segregation, but this implicitly intersects with school and income segregation as well.

2.2.2 Evidence of Residential Segregation

Figure 2 shows the share of each ethnicity in locations across the city per the 2018 census.[8] Whites are concentrated more in the center and outer suburbs of the city. Polynesians predominately reside in South Auckland with an enclave in West Auckland as well. Asians appear more dispersed than Polynesians but with locations of high concentration in the northern, south central, and southeastern suburbs.

[8]As a benchmark, Figure A.14 shows the population distribution in Auckland. Most of the population lives in the central regions of the city, which are largely suburban (88.6% of housing units in Auckland are detached or semi-detached houses).

To quantify the extent of residential segregation, I use the dissimilarity index. This is the mostly widely used measure for segregation. It indicates whether two groups are evenly distributed across the city, with values closer to one (zero) indicating greater (lower) segregation between the two groups. I calculate the Polynesian-white dissimilarity index as 0.545 and the Asian-white dissimilarity index as 0.388. These are moderate levels of segregation per the Massey and Denton (1993) classification. In Appendix A.3, I show other measures of residential segregation (Table A.2) and discuss how Auckland's segregation measures compare to those found in some U.S. cities. The calculations indicate that Auckland exhibits clear residential segregation, and that Polynesians tend to be more segregated than Asians. This is also consistent with the patterns observed in Figure 2.

Figure 2: Ethnic Shares (2018)

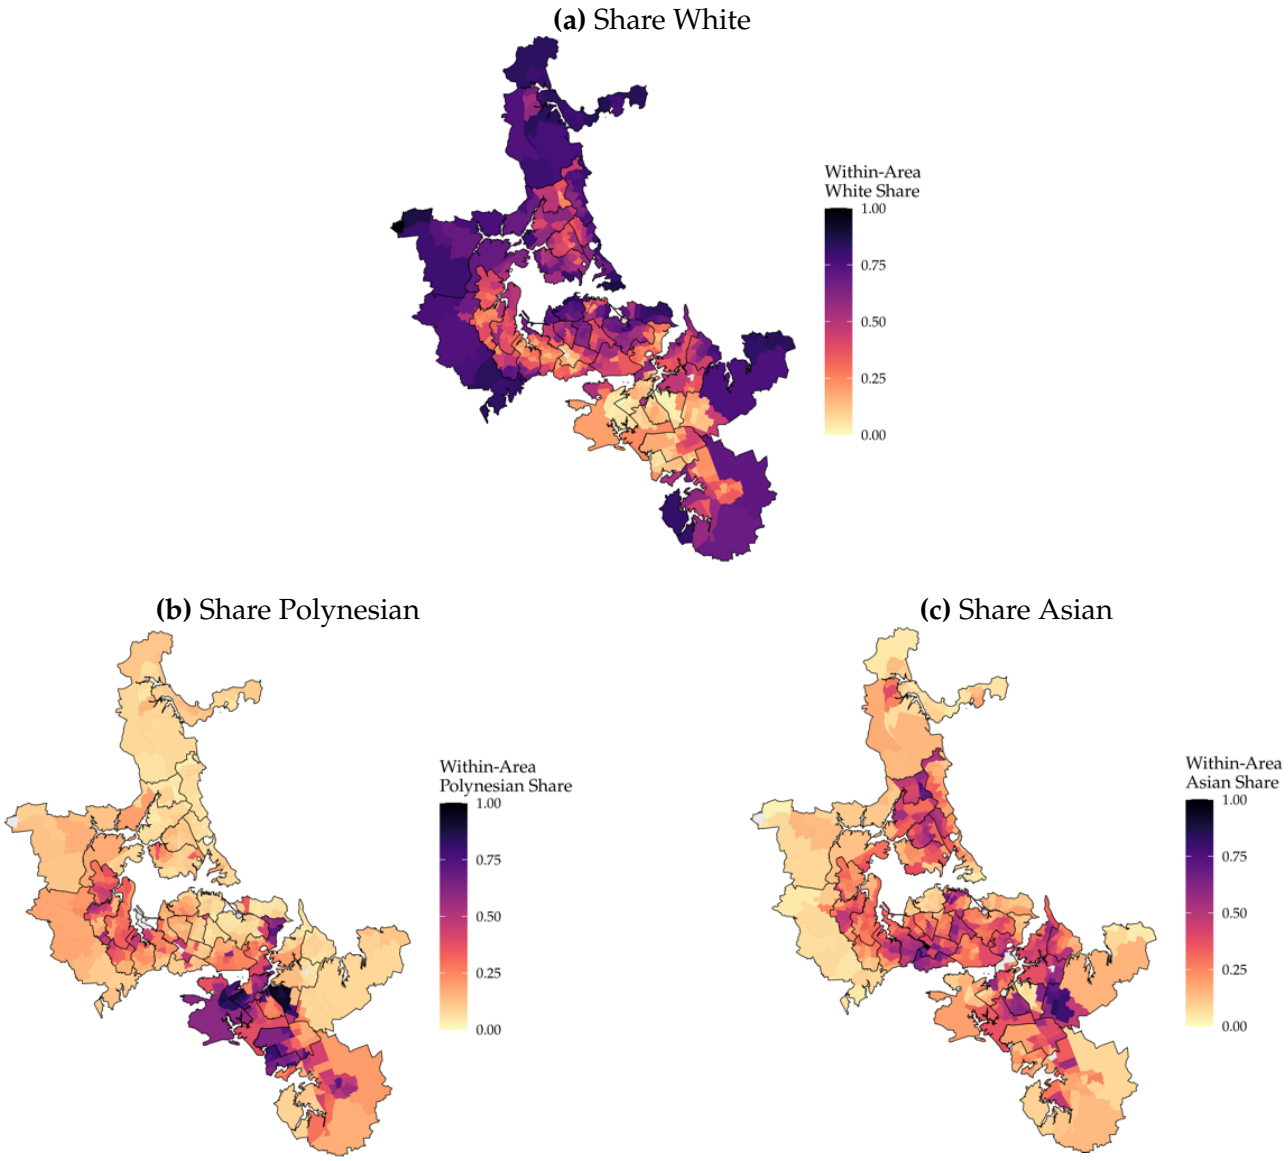

Note: The maps show the share of each location's population belonging to the listed ethnic group. Shares are calculated using 2018 individual-level census microdata. The area units are defined by Statistics New Zealand (Figure A.16b). I overlay the borders of the locations used in the structural model estimation (Figure A.16c).

2.3 Data for Model Estimation

To estimate the structural model, I require panel data on families. This dataset should track information about the family over time, including: the family's home location, the ages/grades of the children, and the school that each child attends. I construct this panel using data from Statistics New Zealand (Stats NZ), which has created a database that links individuals across a variety of administrative records.

To construct the family panel, I rely on four sources of data. First, I use the school records. These are annual files available from 2008 to 2020 that list the enrolled students for each school in the country. Second, I use microdata from the quinquennial census, available for 2013 and 2018. The individual-level responses to the census include details on living arrangements and familial relationships. Third, I use personal demographics data collected by Stats NZ, which includes ethnicity identifications, birthday (month and year), and the identity of the birth parents (via the birth records, if available). Finally, I have a panel of residential locations for each individual in the Stats NZ database. This is constructed by Stats NZ using home address information across all the datasets they collect.

The key steps to constructing the family panel are: linking children to their parents, identifying where the family lives each year, and assigning each family an ethnic group. Appendix A.4 provides details. The family panel covers 2008 to 2019. It contains 302,508 unique families and 3,288,381 family-year observations, including before the family's children are born. Table A.6 provides summary statistics about the panel sample. This shows it is representative of the population of interest (as measured in the census per Tables A.4 and A.5) . This dataset is larger in scope compared to data used in similar studies (see Appendix A.4 for a comparison), but this is necessary to address the high dimensionality of the estimation procedure.

As the panel data relies on a number of sources, I perform a number of validity checks on the data (Appendix A.5). I validate the sample's coverage of Auckland families and find that I capture over 75% of children for most years and age groups. I verify that the family's home location tends to be close to their children's schools, and that this is especially true for primary schools. This is the expected pattern and provides suggestive evidence that the process of linking parents to children and the method of identifying the family home are both likely to be accurate. Finally, I show that the chosen ethnic classifications are appropriate and reflect the ethnic shares of the broader population.

I supplement the family panel with additional data about schools and houses. I use publicly available information about schools (including decile scores) and their school zones from the New Zealand Ministry of Education. From the Auckland Council, I have data on the universe of house sales in Auckland for 2003-2019. This data includes the exact home address, sale price, and house characteristics for each sale. Finally, I have data on the near-universe of rental agreements via Stats NZ. This provides information on the weekly rent amount, property type, number of bedrooms, and the rental start date.[9]

[9]Due to the data structure and confidentiality reasons, I cannot link sale/rental transactions to specific families.

3 Causal and Descriptive Evidence

In this section, I detail key empirical patterns that will motivate the structural model. First, using a boundary discontinuity design, I show that decile scores have a causal effect on house prices. This exercise is useful to show that decile scores can be used in the model as the measure of school quality. The boundary discontinuity design is one of the most widely-used methods in the reduced form literature for estimating the valuation of school quality. Its estimates can then be used as a benchmark for the results of the structural estimation. Second, I provide descriptive evidence using intertemporal moving probabilities to show that a family's moving decisions appear to be affected by their children's schooling. This exercise will also highlight the dynamic factors at play and justify the incorporation of children into a dynamic model of residential choice.

3.1 Valuation of Decile Scores

3.1.1 Identification Strategy

I leverage the school zones in this setting to apply a boundary discontinuity design (BDD). The strategy involves looking at an area close to the boundary of a school zone and comparing houses that are inside the zone ("in-zone") to houses that are outside the zone ("out-of-zone"), as illustrated in Figure 3.

The assumption is that these sets of houses should have similar characteristics (both in terms of house characteristics as well as local neighborhood amenities), except that there is a discontinuous change in expected school quality across the boundary.[10] This means that any observed discontinuous changes in house prices can be attributed to the change in school quality. We can therefore interpret this as a causal estimate of the willingness to pay (WTP) for school quality. The equation to estimate is:

$$\ln p_{hbstv} = \beta_0 + \sum_{\substack{V \in \{\text{Primary,} \\ \text{Intermediate,} \\ \text{High School}\}}} \left(\beta_1^V SQ_{bstv} \cdot \mathbb{1}\{V = v\} \right) + \mathbf{X}_{ht} + \delta_b + \gamma_t + \varepsilon_{hbstv} \tag{1}$$

where the indices refer to house h, on side s of boundary b, in time t, for school level v. Equation 1 is regressing the price of a house (p_{hbstv}) on a measure of school quality at level v for that boundary-side (SQ_{bstv}), as well as house/neighborhood controls ($\mathbf{X}_{ht}$), a set of boundary fixed effects (δ_b) and time fixed effects (γ_t). The identification assumption is that the boundary fixed effects should capture all (unobserved) commonalities on *both sides* of the boundary, which should be all factors except school quality. The coefficient of interest is β_1^V, which captures the willingness to pay for an additional school quality at school level $V \in \{\text{Primary}, \text{Intermediate}, \text{High School}\}$. Most papers that apply the BDD method focus on elementary schools (see Black and Machin (2011) for a recent review). As New Zealand

[10] Here, we can think of "school quality" as a neighborhood amenity available to each house. The method of assigning school quality to each house is discussed in Appendix B.3.

Figure 3: Boundary Discontinuity Design Illustration

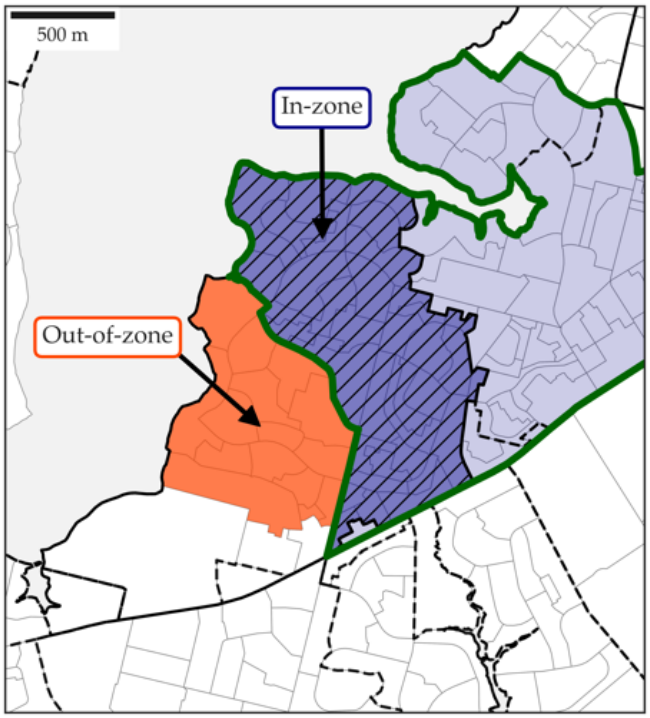

Note: The figure shows an example of the boundary discontinuity design (BDD). The green line indicates the boundary of a one school zone. The blue shaded area is the area inside the zone. The BDD involves comparing an area inside the zone that is close to the boundary (hatched blue) to an adjacent area that is outside of the zone (orange). The dashed black lines indicate the boundary of other school zones.

has school zones for all school levels, I contribute to this literature by testing whether (reduced form) estimates of school quality valuation vary by school level. In Appendix B.1, I show a graphical intuition for the identification strategy. In Appendix B.2, I discuss in detail how boundaries are constructed. In Appendix B.3, I describe how the school quality measure SQ_{bstv} is calculated.

3.1.2 Results

I estimate Equation (1) using a fixed effects regression. I repeat all analysis twice: once using the sales data (Table 1a) and the other using rental data (Table 1b), with column 2 as the main specification . The values shown represent the WTP for an additional decile score as a percentage of the house price.

Using the sales data, I find that high school quality increases house prices by 3.81% for each additional decile score. For primary schools, this value is only 0.87%; I do not find statistically significant results for intermediate schools. I am also able to conduct this analysis using rental data, whereas most papers use sales data. This is a useful comparison as the population who buys houses is likely skewed towards higher income families.[11] The results in Table 1b show that the willingness to pay for high school quality

[11] As Table A.5 shows, a significant portion of families are renters, especially among the Polynesian population.

is 2.8% for an additional decile score. This is approximately three-quarters the magnitude found using the sales data. With the rental data, the estimates for primary and intermediate school quality are not statistically distinguishable from zero, however.

I show additional specifications in Tables 1a and 1b. I include controls for the ethnic shares (column 3), distance to school (column 4), and a running variable to capture unobserved continuous changes in neighborhood quality (column 5). The results hold across these specifications. I conduct further robustness checks in Appendix B.4.

Other papers in the literature find that a 1 standard deviation (s.d.) increase in elementary school quality causes a 3-4% increase in sale prices (Black and Machin, 2011). In Appendix B.4, I argue that 3 decile scores is analogous to a 1 s.d. in school quality. Therefore, I find that a 1 s.d. increase in elementary school quality causes a 2.6% increase in sale prices, but a similar change in high school quality causes an 11.4% increase. This suggests that while the valuation of school quality may be slightly lower in New Zealand as compared to other countries, other papers may be severely underestimating families' full school valuation by only focusing on elementary schools.[12]

Even though a reduced form analysis is possible in this setting, there are two reasons why it is still useful to estimate the willingness to pay for school quality through a structural model.

First, the reduced form estimates are estimated using sales/rental data, without microdata about the families generating these transactions. This is typically the data constraint other researchers face when conducting a BDD analysis. Without microdata, heterogeneity analysis is limited or nonexistent. As we cannot identify which families bought which house, we would not be able to say whether there is variation in quality valuation across, say, ethnic or racial lines. Moreover, families also vary in their composition. This means that the willingness to pay values do not take into account the number of children or their ages. However, these factors influence the "amount" of school quality each family enjoys.[13] The policy-relevant statistic is in fact the willingness to pay for school quality *per child per year* (Caetano, 2019). In contrast, the BDD estimates the willingness to pay for the whole family over their expected tenure in the house.

Second, the reduced form approach can give us estimates for the valuation of school quality, however these alone are not enough to understand how schools affect residential choice. We would need to account for other factors in the family's decision-making process. Then we would also need to understand how the re-sorting of families can have an effect on aggregate market variables (i.e. prices) and the subsequent equilibrium effects. This is beyond the scope of a reduced form framework, but is achievable with a structural model of the market.

[12]The lower WTP estimates here can be explained by the greater degree of school choice in New Zealand. Zheng (2021) embeds an event study into the BDD and finds that the valuation of school quality decreases after charter schools enter.

[13]For example, if every family had two children, then the true willingness to pay for school quality would in fact by half the estimated amount. Similarly, the age of the children determines the horizon over which this school quality will be enjoyed.

Table 1: Boundary Discontinuity Regression Estimates

(a) Sales

	(1)	(2)	(3)	(4)	(5)
School Quality					
High School	0.0709***	0.0381***	0.0334***	0.0350***	0.0407***
	(0.0199)	(0.0122)	(0.0109)	(0.0115)	(0.0129)
Intermediate	0.0187	0.0003	−0.0029	−0.0024	−0.0046
	(0.0279)	(0.0141)	(0.0136)	(0.0137)	(0.0140)
Primary	0.0166*	0.0087***	0.0065**	0.0067**	0.0077**
	(0.0097)	(0.0029)	(0.0032)	(0.0032)	(0.0029)
Other Controls					
House Characteristics	No	Yes	Yes	Yes	Yes
Ethnic Shares	No	No	Yes	Yes	Yes
Distance to School	No	No	No	Yes	Yes
Running Variable	No	No	No	No	Yes
N	128,646	128,646	128,646	128,646	128,646
R^2	0.554	0.813	0.816	0.816	0.816
Clusters	39	39	39	39	39

(b) Rents

	(1)	(2)	(3)	(4)	(5)
School Quality					
High School	0.0323**	0.0283***	0.0270***	0.0281***	0.0259**
	(0.0120)	(0.0097)	(0.0093)	(0.0095)	(0.0097)
Intermediate	0.0171	0.0103	0.0086	0.0086	0.0076
	(0.0106)	(0.0073)	(0.0071)	(0.0068)	(0.0069)
Primary	0.0044	0.0032	0.0002	0.0002	−0.0012
	(0.0046)	(0.0045)	(0.0037)	(0.0038)	(0.0045)
Other Controls					
House Characteristics	No	Yes	Yes	Yes	Yes
Ethnic Shares	No	No	Yes	Yes	Yes
Distance to School	No	No	No	Yes	Yes
Running Variable	No	No	No	No	Yes
N	313,404	313,404	313,404	313,404	313,404
R^2	0.331	0.545	0.551	0.551	0.551
Clusters	39	39	39	39	39

Note: The table shows the estimate of (a) log house sale price and (b) log annual rent on school quality (decile score), boundary and year fixed effects, and the listed coefficients. School quality (decile score) is interacted with an indicator for each school level. In panel (a), house characteristics include the area and age of the house. In panel (b), they are the number of bedrooms and rental type (apartment or house). Ethnic shares are calculated at the family level at each Level 2 area (Figure A.16b). Running variable is distance to boundary $d \in [-500, 500]$, where the side with the higher (lower) average decile score has positive (negative) values. Clusters are set as the Level 3 area units (Figure A.16c). Significance stars thresholds: $^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$. In accordance with Stats NZ confidentiality rules, the number of observations has been randomly rounded to a multiple of 3 (RR3).

3.2 Moving Decisions

The BDD analysis in section 3.1 has shown that parents value school quality (as measured by decile scores). However, it also highlighted the need for a structural model to fully understand the relationship between schools and neighborhood choice. This section will show to what extent schools are affecting the moving probability of families. These descriptive statistics will provide motivation for the dynamic structural model in section 4.

I use the constructed family panel to estimate the probability of moving at different stages for the family. I observe that families are less likely to move as they grow older. Figure B.13 shows that families have an almost 35% probability of moving when the mother is under 25 years old but a less than 10% probability of moving when the mother is over 45. My focus then turns to establishing whether moving patterns change differentially over time relative to the child's age and grade.

Figure 4 shows how the family's probability of moving varies by the child's grade (where a negative grade indicates the years before the child starts school). For ease of interpretation, I only focus on families with one child and control for mother's age to account for the downward trend observed in Figure B.13.[14] The figure shows that there are local peaks in moving probabilities when the child is born (grades –5 or –4), and when they start primary school (grade 1) or high school (grade 9).[15]

While these are not causal results, this is a stark finding as moving probabilities tend to change smoothly over time – for example as shown in Figure B.13 (relative to mother's age) and Figure C.7 (relative to housing tenure). However, the birth of children and the changes in school levels generates discontinuities in the needs of families. This is reflected by the sudden spikes in the intertemporal moving profile at those specific points in time. This suggests that children and schools are indeed playing a role in the family's moving decisions.

The empirical patterns also suggest that families exhibit dynamic behavior in terms of how children affect their moving decisions. If families were myopic, we would expect to see much larger spikes at the start of each school level (e.g. in such a world, families would not anticipate that their child would enter primary school and so would have to re-adjust suddenly once the child begins school). In fact, the reduced form estimates suggest that families value high school more than primary school. Yet parents are more likely to move when their child enters primary school than when they enter high school. This can be rationalized if moving costs are high and families are forward-looking about future benefits when considering where to live (Bayer et al., 2016).

[14] For every family-year observation in the panel, I regress a moving indicator on dummies for the grade of the child and dummies for the age of the mother. The probabilities plotted in the figure are the coefficient estimates on dummies for each child's grade added to the coefficient for an age 35 mother. I choose 35 as a baseline since this is the average age of mothers at the time of their child's birth.

[15] Children in New Zealand are either age 5 or 6 in the first grade. Therefore, the grade values –5 and –4 can both represent the period in which the children is born. The majority of children are age 6 in first grade (i.e. are born in grade –5).

Figure 4: Moving Probabilities – Child's Grade

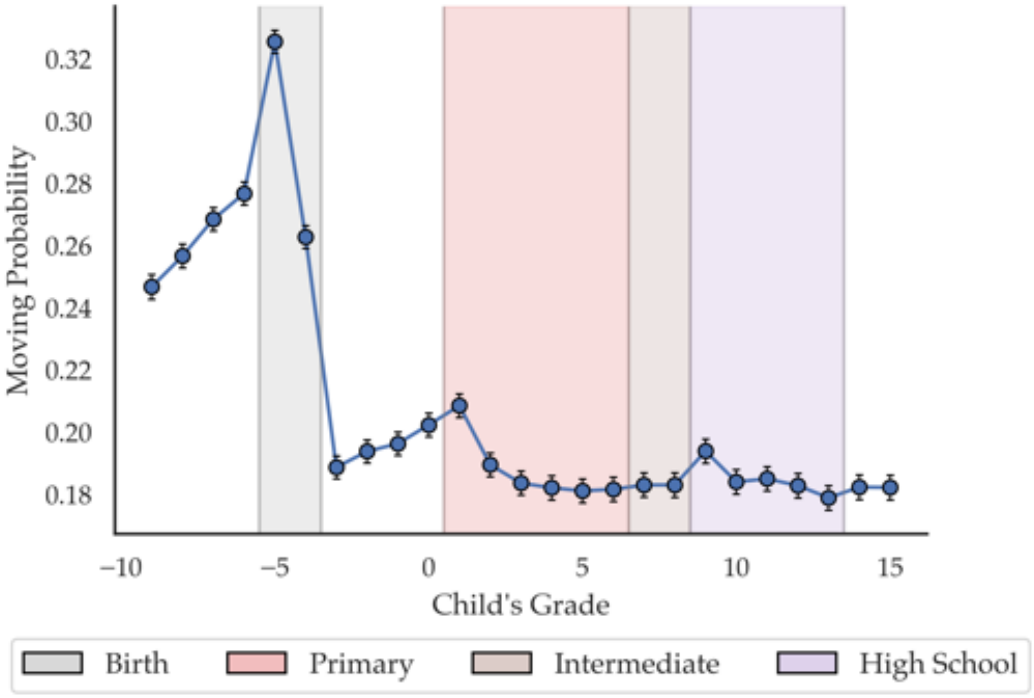

Note: The figure shows the moving probability given the grade of the child. I regress a moving indicator on dummies for each grade of the child and dummies for each age of the mother. Each point is equal to the grade-specific coefficient plus the coefficient for age 35 mother. 95% confidence intervals are plotted. Depending on the age they start school, the child's birth is represented at either grade −5 or −4. Shaded regions indicate the ranges for each school level: primary (grades 1-6), intermediate (grades 7-8), and high school (grades 9-13).

In Appendix B.5, I show further analysis of the intertemporal moving profile. I show that the moving profile varies by the birth order of the child (Figure B.14). For example, families are far less likely to move at the time of youngest child's birth as compared to the time of the eldest child's birth. This implies that a structural model needs to take into account different family sizes and the children's birth order when considering the family's moving choice. I also validate that schools (and not other factors) are driving the effects seen in Figure 4. I leverage cut-offs in school entry rules to identify plausibly exogenous variation in the child's grade. I compare children who were born in the same calendar year but are assigned to different grades (Figure B.16). This shows that after a child enters school, it is their grade (rather their age) that is affecting the the family's moving decisions.

Finally, it is important to note that across all these figures, we see that most of the family's moves occur before the child even begins school (and oftentimes even before their birth). This highlights the data challenge of studying this topic as it requires tracking the family's home location over a long period of time. For example, using home addresses from a child's education records would be insufficient for studying residential sorting. This data would not capture the family's location in the years prior to the child entering school. This is a challenge that I overcome by using the New Zealand administrative data.

4 Model of Residential Choice

In this section, I develop a residential choice model using a dynamic discrete choice framework. This model studies the within-city neighborhood choice of households, similar to Almagro and Domínguez-Iino (2021). The main contribution to the literature is that the model incorporates children into the agent's problem.[16] Following Kalouptsidi et al. (2020), I then show how the model implies a linear equation from which the key parameters can be estimated.

4.1 Setup

In the model, each agent i is a family. This represents a household consisting of adults who currently have (or will have) one or more children.[17] Each family is of type $\psi \in \Psi$ (this will represent ethnic groups in the estimation). Time is discrete and indexed by $t \in \{1, ..., T\}$. This corresponds to one calendar year.[18] In the city, there are L locations, indexed by $l \in \mathcal{L} = \{1, ..., L\}$. An additional location $l = 0$ will represent the outside option, which represents all areas outside of the city. The full set of home locations is therefore $\mathcal{L} \cup \{0\}$.

Each family i has $C_i \in \mathbb{N}^+$ children. A child $c \in C_i$ is born in period T_{ic}^b and therefore is of age $a_{itc} = t - T_{ic}^b$ at time t. In Appendix C.1.1, I outline the process in which children are born and families enter into the economy. The key element is that while I make the number of children (C_i) known and exogenous, I introduce randomness in the timing of each child's birth. Families anticipate the birth of their children by a fixed number of periods (denoted as $\underline{b}$). Let C_{it} be the number of children in family i as of time t.

4.2 The Household's Problem

4.2.1 Decisions

Each period, the family decides whether to stay in their current home or move to a new neighborhood. Call the location decision as $d_{it} \in \mathcal{D} = \mathcal{L} \cup \{0, stay\}$, where $d_{it} = stay$ corresponds to staying in their current home and $d_{it} = 0$ corresponds to leaving the city (outside option). Note that for a family living in location l, a choice of $d_{it} = l$ means moving to a different house in the same area. This is a different choice to $d_{it} = stay$, which means not moving.

[16] Caetano (2019) is most similar to this paper in its incorporation of children. However, the key difference between our papers is that I model each family individually and fully incorporate multiple children into the family payoff. In contrast, he categorizes families based on the age of their eldest child and accounts for other children based only on observed group averages. Our differences in data further change our modeling approaches. I have panel data on all families, allowing me to specify how the family composition transitions. His data is cross-sectional and so he makes a synthetic cohort assumption for estimating the transitions, which also circumvents him having to model the transitions explicitly.

[17] The number of adults is irrelevant in the model. I abstract away from changes in parents due to events such as marriages, divorces, separations, and deaths. I also do not focus on families who will never have children over their lifetime.

[18] Fortunately, as New Zealand is in the Southern Hemisphere, the academic year aligns with the calendar year.

4.2.2 Maximization Problem

The family's problem is to choose $\{d_{it}\}_{t=T_i^0}^{T_i^1}$ to maximize their expected discounted sum of payoffs:

$$E\left[\sum_{t=T_i^0}^{T_i^1} \beta^{t-T_i^0} \Pi\left(d_{it}, \mathbf{s}_{it}\right) \middle| \mathcal{I}_{iT_i^0}\right] \tag{2}$$

where T_i^0 and T_i^1 (defined in section C.1.1) represent $\underline{b}$ periods before the eldest child's birth to the adulthood of the youngest child. β is the discount rate and $\mathcal{I}_{iT_i^0}$ is the family's information set at T_i^0. Π is the payoff in each period. This depends on the decision d_{it} and the state variables $\mathbf{s}_{it} = (k_{it}, w_t, \eta_t, \varepsilon_{it})$. Following Kalouptsidi et al. (2020), the state variables are grouped as follows:

- k_{it} are observable individual-level variables, which includes: the family's type (ψ_i); current home location ($h_{i,t-1}$); the location tenure ($\tau_{i,t-1}$), defined as number of years lived in current home; the current number of children ($C_{i,t-1}$); as well as the current age ($a_{i,t-1,c}$), grade ($g_{i,t-1,c}$), and school ($m_{i,t-1,c}$) for each child $c \in C_{i,t-1}$.[19]
- w_t are observable neighborhood characteristics, which includes: school-related variables (e.g. school quality); house prices (r_t) in each area; and the share of each type ψ in the city's locations (TS_t^{ψ}).
- η_t are unobservable neighborhood characteristics of each location.
- $\varepsilon_{it} = (\varepsilon_{it0}, \varepsilon_{it1}, ..., \varepsilon_{itL}, \varepsilon_{it,stay})$ are unobservable idiosyncratic action-specific shocks.

Note that k_{it} is a vector of $t-1$ variables, such as home location ($h_{i,t-1}$). These $t-1$ variables represent the family's *current* living situation at the start of period t and are simply expressed as k_{it} for notational convenience. In Appendix C.1.2, I specify the intraperiod timing and what families know when making their decision.

4.3 Per Period Utility

The per period payoff function Π for a decision $d_{it} = j \in \mathcal{D}$, given the state variables $\mathbf{s}_{it}$, is defined as:

$$\begin{aligned} \Pi(j, \mathbf{s}_{it}) &= u(j, k_{it}, \omega_t) + \varepsilon_{itj} \\ &= \bar{u}(j, k_{it}, w_t) + \xi(j, \eta_{it}; k_{it}, w_t) + \varepsilon_{itj} \end{aligned} \tag{3}$$

$u(j, k_{it}, \omega_t)$ is the payoff before the idiosyncratic shock is realized. This can be further separated into $\bar{u}$, utility that is a function of the observable state variables (k_{it}, w_{it}), and ξ, utility that is a function of the

[19] Children of the same age can be in different grades (Appendix B.5), which requires specifying them as separate variables.

unobservable market-level state variables (η_t).[20] Equation (3) assumes that the payoffs are additively separable in both the unobservable neighborhood payoffs ξ and the idiosyncratic error ε_{itj}.

Next, I will outline the components of the observable utility $\bar{u}$, whose parameters will be estimated in section 5. For brevity, I will provide descriptions of each component and leave the full mathematical expressions to Appendix C.1.3. I provide a summary of the observable utility components in Table C.1.

4.3.1 School Variables

A family's location decision affects what school their child attends and therefore their payoff from the school variables. I will outline how children are assigned to schools in section 4.4. In this part, I consider the payoff associated with schools conditional on the child's grade (g_{it}), school (m_{it}), and family home location (h_{it}). For ease of exposition, I omit the c subscript and momentarily assume the family only has one child. For families with multiple children, I sum the payoffs associated with each child. There are four parts to the school payoff: school quality, a school moving cost, distance to school, and an outside option value.

The school quality component is denoted by $SQ(k_{i,t+1}, w_t)$. This represents the quality of the child's school. The valuation for quality is determined by a coefficient θ_ψ^V, which varies by the child's school level $V \in \{\text{Primary}, \text{Intermediate}, \text{High School}\}$ and the family's type $\psi \in \Psi$.

The school moving cost is denoted by $SMC\left(k_{i,t+1}, m_{i,t-1}\right)$. This represents the cost of changing a child's school within a given level. This captures, for example, the inconvenience of changing schools or the psychological cost of the child losing connections to their friends and teachers. I set this cost as υ_ψ^V, which (like the quality valuation) varies by school level V and family type ψ. This cost is experienced when the child's school differs from their previous year's school ($m_{it} \neq m_{i,t-1}$). Importantly, there is no school moving cost when a child begins a new school level (as they always have to change schools when this happens). This occurs at $g_{it} = 1, 7$, and 9, which are the starting grades for primary, intermediate, and high school, respectively.[21]

Finally, the distance to school and outside options components also have payoffs that differ by school level and family type. Distance is denoted by $DS(k_{i,t+1})$, and it represents the distance from the family's home location to the child's assigned school. The outside option is $OOS(k_{i,t+1})$, which reflects that parents may send their children to private schools within the city as well as schools outside of the city. Further details of its parameterization are left to Appendix C.1.3.

[20] I separate out the inputs of ξ to make it clear that it should be interpreted as the utility from unobservable neighborhood characteristics. ξ is capturing anything that is not explained by $\bar{u}$. ξ is written as a function of the observable states (k_{it}, w_t) as they may play a role in determining the utility from η_t. For example, if the action chosen is $j = stay$, then knowing the previous home location $h_{i,t-1}$ (part of k_{it}) is necessary to knowing the family's location in the next period (and therefore the unobserved market-level utility they experience in that location).

[21] I parameterize the transition to a new school level as costless because there are no opportunity costs to be captured. The school moving cost arising from changing school levels will be incurred regardless of the family's action.

Per the notation in section 4.2, the grade (g_{it}), school (m_{it}), and home (h_{it}) are part of $k_{i,t+1}$. This is unknown at the time families choose their action $d_{it} = j$. Therefore, the school payoff component is expressed as an expectation over $k_{i,t+1}$:

$$E_k\Big[SQ(k_{i,t+1}, w_t) + DS(k_{i,t+1}) + SMC(k_{i,t+1}, m_{i,t-1}) + OOS(k_{i,t+1})\Big|j, k_{it}, w_t\Big] \tag{4}$$

In section 4.4, I will provide the relevant probability distributions to show how this expectation is evaluated. Finally, note that if a child is not in school, then the school payoffs are all equal to zero.

4.3.2 Neighborhood Variables

Let $r(l)$ be the average house price in location l and $TS^{\varphi}(l)$ be the share of location l's families that are of type $\varphi \in \Psi$. The neighborhood component of utility is given by:

$$E_k\left[\alpha^r_{\psi} \ln r\left(h_{it}\right) + \sum_{\varphi \in \Psi} \alpha^{\varphi}_{\psi} TS^{\varphi}\left(h_{it}\right)\Big|j, k_{it}, w_t\right] \tag{5}$$

where α^r_{ψ} captures the price sensitivity of group ψ and α^{φ}_{ψ} captures the preferences of type ψ to the share of type φ living in their neighborhood. Under this specification, all types experience the same house prices in a given location. This implicitly assumes away potential frictions that some groups may face in the housing market (e.g. differential bargaining power or discrimination) as well as within-location sorting (wherein some types sort into lower value housing even within a neighborhood).

4.3.3 Moving and Location Variables

The final set of payoffs relate to moving and the family's home location:

$$E_k\Big[\delta_{\psi}\left(h_{it}\right) + \alpha^{\tau}_{\psi}\tau_{it}\Big|j, k_{it}\Big] + MC\left(j, k_{it}\right) \tag{6}$$

First, there is a location fixed effect, which is $\delta_{\psi}(l)$ for location l. This captures time-invariant characteristics of locations over the sample period. Second, there is a tenure variable τ_{it}, which is the number of years the family has been in their home. Finally, there is a moving cost denoted by $MC\left(j, k_{it}\right)$, which I parameterize as a fixed cost plus a per kilometer distance cost. Moves between the city and the outside option are equal to a different fixed cost. Families who stay in their current home do not pay a moving cost. I allow preferences for all these components to vary by family type.

These modeling choices are supported by the data. Figure C.7 shows that the probability of moving is decreasing relative to the tenure length. Tenure could act as a proxy for connections to the current house and the local area (e.g. knowing your neighbors). Figure C.8 shows the distribution of moving distances: within-city moves are relatively close with the median distance being 3.8 kilometers. This supports parameterizing moving costs as a function of distance.

Putting everything together (expressions 4, 5, 6), the observable payoff for an action $d_{it} = j$ is:

$$\begin{aligned}\bar{u}(j, k_{it}, w_t) = E_k \Big[& SQ(k_{i,t+1}, w_t) + DS(k_{i,t+1}) + SMC(k_{i,t+1}, m_{i,t-1}) + OOS(k_{i,t+1}) \\ & + \alpha^r_\psi \ln r\left(h_{it}\right) + \sum_{\varphi \in \Psi} \alpha^\varphi_\psi TS^\varphi\left(h_{it}\right) \\ & + \delta_\psi\left(h_{it}\right) + \alpha^\tau_\psi \tau_{it} \Big| j, k_{it}, w_t \Big] + MC\left(j, k_{it}\right) \end{aligned} \tag{7}$$

4.4 State Transition

4.4.1 Individual State Transition

Equation (7) shows that the observable utility is a conditional expectation over how k_{it} transitions to $k_{i,t+1}$. In this section, I will outline this transition process. While some variables transition deterministically (making the expectation trivial), the child-related variables follow a stochastic process. I first make the following assumption about the individual state transition:

Assumption 1: *The transition of the observable individual state variable k_{it}, denoted by F^k, depends only on the agent's action and the observable states:*

$$F^k\left(k_{i,t+1} \middle| d_{it}, k_{it}, w_t\right)$$

Next, the home location updates according to where the family chooses to move, or remains the same if they do not move:

$$h_{it} = \begin{cases} h_{i,t-1} & \text{if } d_{it} = stay \\ d_{it} & \text{if } d_{it} \in \mathcal{L} \cup \{0\} \end{cases}$$

Tenure accumulates as long as the family stays in the same home (up to some maximum value, $\bar{\tau}$). If the family moves to a new home (even in the same location), the tenure counter resets to zero.[22]

$$\tau_{it} = \begin{cases} \min\left\{\tau_{i,t-1} + 1, \bar{\tau}\right\} & \text{if } d_{it} = stay \\ 0 & \text{otherwise} \end{cases}$$

[22] An argument could be made that tenure should continue to accumulate if households stay in the same location ($d_{it} = h_{i,t-1}$). This depends on how large the locations are defined and whether the neighborhood connections that τ captures are highly local to the house or represent a broad area. If tenure captured these broader neighborhood connections, we should see that higher tenure implies closer moves. Figure C.9 shows a negative relationship between moving distance and tenure only for Polynesians (that eventually plateaus), while a weaker correlation for the other two groups.

The age of each child c grows by one each year:

$$a_{itc} = a_{i,t-1,c} + 1, \forall c \in \{1, ..., C_{it}\}$$

The child's school (m_{itc}) is determined by a stochastic school assignment process. For this, I define a set of locations $\mathcal{R}(j) \subset \mathcal{L}$ that are nearby to home location j. I call this the region of j, with $j \in \mathcal{R}(j)$. If a child enters a new school level, i.e. $g_{it} \in \{1, 7, 9\}$, they change schools with probability 1. If they are in the same school level as before and their family continues to live close to their previous location, i.e. $h_{it} \in \mathcal{R}(h_{i,t-1})$, then the child changes school with probability ρ. However, if the family moves to a location that is far from their original home, i.e. $h_{it} \notin \mathcal{R}(h_{i,t-1})$, then their child changes school with probability 1. I refer to the parameter ρ as the re-assignment probability. When a child changes school, they receive a draw from the probability distribution $M_t^{New}(h_{it}, g_{itc})$. This randomly assigns a child to a school based on the family's home location, the child's grade, and the time t. This process is summarized in Figure 5.

Figure 5: School Assignment Process

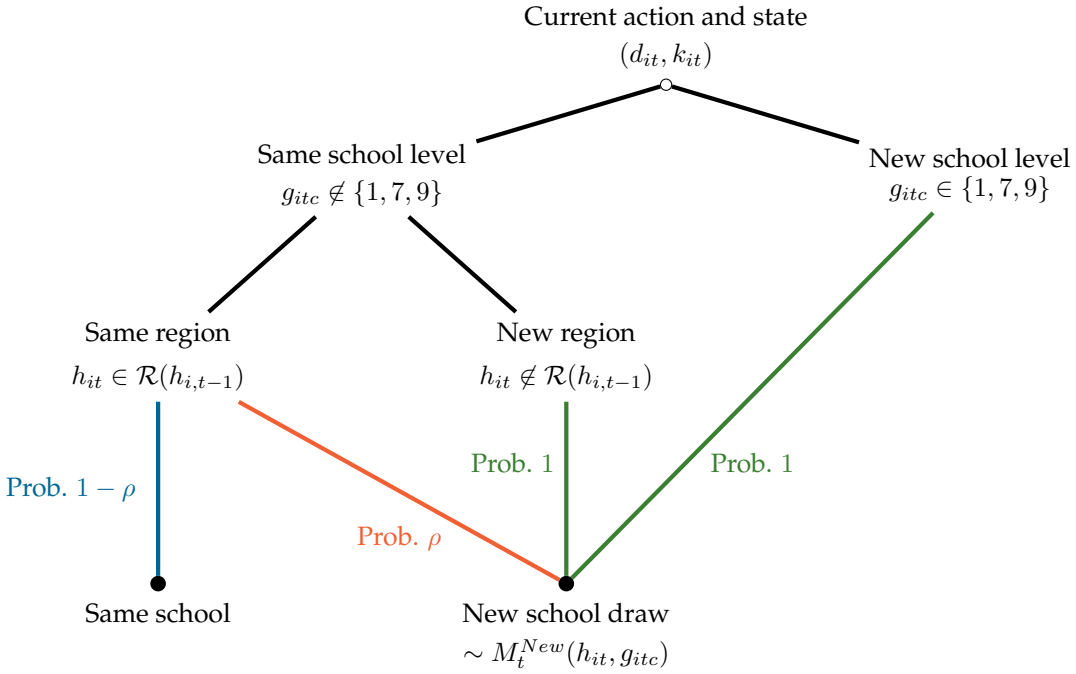

In Appendix C.1.4, I give further details on the school assignment process and show that it is supported by patterns in the data. I also outline the transition process for the other child-related variables, namely the number of children (C_{it}) and the grade of the child (g_{itc}).

4.4.2 Full State Transition

Next, I describe the transition process for the remainder of the state variable $\mathbf{s}_{it} = (k_{it}, w_t, \eta_t, \varepsilon_{it})$:

Assumption 2: *The transition of the market-level states $\omega_t = (w_t, \eta_t)$ is given by: $F^{\omega}(\omega_{t+1}|\,\omega_t)$.*

Assumption 3: *The ε_{it} shocks are i.i.d. and follow a Type 1 extreme value distribution.*

Assumption 2 implies that individual agents are small relative to the market. Therefore, no individual agent's action affects the market-level variables. This is appropriate here as an agent is one family, while the market is an entire city. Assumption 3 is standard in the industrial organization literature. Assumptions 1, 2, and 3 imply that the state variable $\mathbf{s}_{it}$ follows a first order Markov process with a transition probability given by:

$$F(\mathbf{s}_{i,t+1}|\,d_{it}, \mathbf{s}_{it}) = F^{k}(k_{i,t+1}|\,d_{it}, k_{it}, w_t)\, F^{\omega}(\omega_{t+1}|\,\omega_t)\, F^{\varepsilon}(\varepsilon_{i,t+1}) \tag{8}$$

4.5 Deriving the Estimation Equation

In this section, I apply the Euler Equations in Conditional Choice Probability (ECCP) estimator proposed by Kalouptsidi et al. (2020). In Appendix C.3, I provide an intuition of the ECCP estimator through a simplified example. In Appendix C.4, I show the full derivation of the main equation.

Applying ECCP relies on two features of the model. The first feature is that agents (families) are small compared to the market (the city), i.e. Assumption 2. This allows me to use observed values of market variables, rather than their expected value. This means I do not need to estimate the transition process of the market variables, nor do I need to specify how agents *believe* the market variables will evolve. The second feature is that there is finite dependence in the agent-specific variables (Arcidiacono and Miller, 2011). By applying this property, I can compare the payoff from two action sequences such that continuation values get cancelled out. This circumvents having to calculate the full value functions, and thus dramatically simplifies the computations.

4.5.1 Renewal Actions

Let $\mathcal{K} = \{k^1, ..., k^K\}$ represent the set of possible values for k_{it}. For a given state (k_{it}, ω_t), we want to consider two distinct actions d and d' taken in period t, followed by an action d'' in period $t+1$. For the remainder of the derivations, I hold the individual state variable fixed at $k_{it} = k^*$.

Let $\mathbf{F}^k_{jt}$ be the $K \times K$ transition matrix of k_{it} given an action j and market state w_t. Let $F^k_{jt}(k^*)$ be a $1 \times K$ vector equal to the row in $\mathbf{F}^k_{jt}$ corresponding to $k_{it} = k^*$.[23] The action d'' is chosen such that the following condition holds:

$$F^k_{dt}(k^*)\mathbf{F}^k_{d'',t+1} = F^k_{d't}(k^*)\mathbf{F}^k_{d'',t+1} \tag{9}$$

In other words, choosing $d_{it} = d$ followed by $d_{i,t+1} = d''$ must result in the same distribution of states

[23] The (a, b) entry of $\mathbf{F}^k_{jt}$ is equal to the probability $F^k\left(k_{i,t+1} = k^b \middle|\, d_{it} = j, k_{it} = k^a, w_t\right)$. $F^k_{jt}(k^a)$ is the a^{th} row of $\mathbf{F}^k_{jt}$.

(in period $t+2$) if the family had instead chosen $d_{it} = d'$ in period t. This is a case of one-period finite dependence (Arcidiacono and Miller, 2011). Note that d'' is chosen depending on k^*, d, and d'. I will refer to d'' as a *renewal action*.

To ensure Equation (9) holds, a renewal action d'' must result in: (i) the tenure capital τ being reset to zero, and (ii) the children's school assignment being re-drawn from the school distribution M_t^{New}. Therefore, a renewal action must be a move to a new region. Specifically, let h_{it} be the family's location after action d and h'_{it} be their location after action d'. The condition for d'' to be a renewal action is then:

$$d'' \in \mathcal{D} \backslash \left(\{stay\} \cup \mathcal{R}\left(h_{it}\right) \cup \mathcal{R}\left(h'_{it}\right) \right) \tag{10}$$

4.5.2 Equations

I define $p_{jt}(k^*)$ as the conditional choice probability (CCP) of an action $d_{it} = j$ given the state $k_{it} = k^*$.

$$p_{jt}(k^*) \coloneqq Pr\left(d_{it} = j \middle| k_{it} = k^*, \omega_t\right)$$

Let $\mathbf{p}_{jt}$ be the $K \times 1$ vector of CCPs: $\mathbf{p}_{jt} = [p_{jt}(k_1), ..., p_{jt}(k_K)]'$. Taking the transition of ω_t as given, and starting at $k_{it} = k^*$, the log likelihood of choosing the action sequence $(d_{it}, d_{i,t+1}) = (j, d'')$ is:

$$Pr\Big((d_{it}, d_{i,t+1}) = (j, d'') \Big| k_{it} = k^*, \omega_t\Big) = \ln p_{jt}(k^*) + \beta F_{jt}^k(k^*) \ln \mathbf{p}_{d'',t+1} \tag{11}$$

The second term of expression (11) represents the likelihood of choosing $d_{i,t+1} = d''$, while taking into account that choosing $d_{it} = j$ affects the state transition of $k_{i,t+1}$ (starting from $k_{it} = k^*$).

We can then evaluate expression (11) for the two action sequences: $(d_{it}, d_{i,t+1}) = (d, d'')$ and (d', d''). Taking the difference between these evaluations defines their relative likelihood:

$$Y_{d,d',d'',t}(k^*) \coloneqq \ln\left(\frac{p_{dt}(k^*)}{p_{d't}(k^*)}\right) + \beta\left(F_{dt}^k(k^*) - F_{d't}^k(k^*)\right) \ln \mathbf{p}_{d'',t+1} \tag{12}$$

$Y_{d,d',d'',t}(k^*)$ is the relative likelihood of observing (d, d'') as compared to observing (d', d''), starting at $k_{it} = k^*$ in time t. The model implies that $Y_{d,d',d'',t}$ is explained by the difference in the payoffs between the action sequences (d, d'') and (d', d''). There are two components to this difference.

In period t, there is the difference in the flow utility (exclusive of ε_{it}) of actions d and d'. This is represented by $u\left(d, k^*, \omega_t\right) - u\left(d', k^*, \omega_t\right)$. To simplify notation, I will re-write the previously defined functions with a j and t subscript to represent an input of $d_{it} = j$ and ω_t, respectively. For brevity, I will then omit the function's inputs (including $k_{it} = k^*$, as that is being held fixed). This means the flow utility difference is given by $u_{dt} - u_{d't}$. By Equation (3), this can be decomposed into: $(\bar{u}_{dt} - \bar{u}_{d't}) + (\xi_{dt} - \xi_{d't})$.

In period $t+1$, given that d'' is a renewal action, the realized continuation values are the same regardless of the past actions. Consequently, after families pay a cost to move to location d'', there is no difference in expected payoffs between the two action sequences (given the transition of ω_t). The costs entail the school moving costs (SMC) and the location moving costs (MC).[24] I denote this difference by: $\left(SMC_{d,d'',t} - SMC_{d,d'',t} + MC_{d,d''} - MC_{d',d''}\right)$.[25] However, note that these costs are experienced one period in the future. This means they should be discounted by β. More importantly, it also means that the payoffs depend on how agents forecast the transition process of ω_t. The ECCP estimator involves decomposing the expected continuation valuation into the realized component and an expectational error component. This step is made possible by Assumption 2.

This intuition implies the following relationship:

$$Y_{d,d',d'',t} = (\bar{u}_{dt} - \bar{u}_{d't}) + \beta\left(SMC_{d,d'',t} - SMC_{d,d'',t} + MC_{d,d''} - MC_{d',d''}\right) + \epsilon_{dd't} \tag{13}$$

The error term ($\epsilon_{dd't}$) has two parts. The first is the difference in unobserved neighborhood utility after taking each action ($\xi_{dt} - \xi_{d't}$). The second is the difference in expectational error of the continuation value after action d as compared to after d'.

Equation (13) can be re-written using the parameterization of $\bar{u}_{dt}$ in Equation (7). This gives the main estimation equation:

$$\begin{aligned} Y_{d,d',d'',t} = &(SQ_{dt} - SQ_{d't}) + (DS_d - DS_{d'}) + (SMC_{dt} - SMC_{d't}) \\ &+ \alpha^r_\psi (\ln r_{dt} - \ln r_{d't}) + \sum_{\varphi\in\Psi} \alpha^\varphi_\psi \left(TS^\varphi_{dt} - TS^\varphi_{d't}\right) \\ &+ \left(\delta_{\psi d} - \delta_{\psi d'}\right) + \alpha^\tau_\psi (\tau_d - \tau_{d'}) + (MC_d - MC_{d'}) \\ &+ \beta\left(SMC_{d,d'',t} - SMC_{d,d'',t} + MC_{d,d''} - MC_{d',d''}\right) + \epsilon_{dd't} \end{aligned} \tag{14}$$

We can take Equation (14) to the data as the left hand side is a function of empirical probabilities. The right hand side is a function of data values, state transition probabilities, and the model parameters (which are to be estimated). Implicitly, this equation depends on the choice of the individual state variable $k_{it} = k^*$ and the time t (which captures the variation in market variables ω_t). The set of renewal actions d'' is determined by k^*, d, and d'. Therefore, the data used to estimate Equation (14) should be at the (k^*, t, d, d', d'') unit level.

[24] For (d, d''), a family has to move from h_{it} to d''. For (d', d''), the family has to move from h'_{it} to d''. This generates the difference in MC. The SMC difference depends on the likelihood that the child changes school after moving to d''. For the most part, SMC will be identical regardless of the d_t action (d or d'), which means it gets differenced out completely.

[25] Concretely, $SMC_{j,d'',t} + MC_{j,d''} = E_k\left[SMC(k_{i,t+2}, m_{it}) + MC\left(d'', k_{i,t+1}\right) | \, d_{i,t+1} = d'', d_{it} = j, k_{it} = k^*\right]$ for $j = d, d'$.

5 Estimation

In this section, I outline the estimation of the structural model parameters. I first outline how the data is mapped to the model. This involves estimating conditional choice probabilities, estimating state transitions, and calculating the payoff components for state-action tuples. Then, I focus on estimating the parameters of Equation (14).

The key empirical challenge for the estimation is accounting for endogeneity. It is likely that components of the observable utility $\bar{u}_{dt}$ are correlated with components of the unobservable utility ξ_{dt}. For example, areas with higher school quality are also likely to have other desirable neighborhood characteristics. As ξ_{dt} is part of the error term in Equation (14), this means an OLS regression will provide biased estimates of the parameters.

To account for this, I estimate the parameters in a three-step procedure. First, to estimate the parameters on the school variables, I difference out the unobserved neighborhood characteristics by comparing families *with* school-aged children to otherwise similar families *without* school-aged children. Second, I use an instrumental variable regression to identify the parameters for price and ethnic shares. I use a Hausman-style instrument of non-neighboring prices and a shift-share instrument that simulates ethnic shares based on past settlement. Finally, I estimate the remaining parameters with OLS as there is limited scope for endogeneity once the neighborhood variables are accounted for.

5.1 Setting the State Space

I provide an overview of how I define the state space of (k_{it}, ω_t), with further details outlined in Appendix D.1. Each period t is one year. The data covers 12 periods, from 2008 to 2019 (inclusive). Calendar years represent the market-level state ω_t. For the individual-level state, the family types are set as ethnicities: $\Psi = \{\text{White, Polynesian, Asian}\}$. I set the maximum tenure value as $\bar{\tau} = 5$. To keep the state space manageable, I restrict to families with at most 2 children. I also restrict the decision window to be from 3 years prior to the eldest child's birth up to when the youngest child reaches 25. These choices mean that I observe 674 possible child-age-grade combinations.

For the locations, I discretize the city area as shown in Figure A.16. I set the locations as the Level 3 areas (Figure A.16c). The mean number of families per location is around 5,800 over the sample period. There are $|\mathcal{L}| = 42$ locations in the city. This means there are $|\mathcal{D}| = 2 + |\mathcal{L}| = 44$ possible actions (each location plus the stay and outside option actions).

With this setup, the possible number of states is equal to: 12 (the number of years) $\times$ 3 (the number of types) $\times$ 43 (the number of home locations) $\times$ 674 (the number of child-age-grade combinations) $\times$ 6 (the number of tenure values) $=$ 6,260,112.

One component of the state variables that I did not include is the child's school. With over 300 schools in the city, this would greatly increase the already large state space and make the estimation infeasible. However, given that schools are assigned in an exogenous manner (Equation C.7 in Appendix C.1.4), I interpret the variables as weighted sums, with the school attendance probabilities as the weights. In Appendix D.4, I discuss what this means for the values of the school-related variables.

5.2 Estimating Probabilities

I estimate the conditional choice probabilities (CCPs) p_{dt}, which are used for constructing the relative likelihood $Y_{d,d',d'',t}$ (Equation 12). However, the state space dimension and the relatively large number of actions mean that many state-action combinations have few or no observations.

I apply the semi-parametric estimation method of Raval et al. (2017) to reduce the dimensionality of the problem. This defines an algorithm wherein states with similar characteristics are iteratively grouped together. Within this group, CCPs are estimated using standard techniques. For example, one grouping could be families who are similar on all dimensions, except their tenure capital can be either $\tau = 3$ or 4. I discuss the grouping process further in Appendix D.2.1.

I use the grouping algorithm within a two-step CCP estimation procedure. In the first step, I estimate the conditional probability of choosing between the actions $\{$*stay*, *move*, *outside*$\}$, where *move* indicates all choices $d \in \mathcal{L}$. Reducing the number of actions generates more precise frequency estimates. In the second step, I subset to families who chose *move* in the first step and use a spatial kernel density estimator to calculate choice probabilities. I combine the probabilities from the two steps to derive CCPs for each $d \in \mathcal{D}$. Appendix D.2.2 provides further details on this procedure.

I estimate individual state transition function F^k_{dt} using empirical frequencies. This is used for calculating $Y_{d,d',d'',t}$ (Equation 12) and evaluating the expectations in $\bar{u}_{dt}$ (Equation 7). Appendix D.3 provides further details.

5.3 Components of the Estimating Equation

Equation (14) is at the state-action tuple level (k^*, t, d, d', d''). The data used in the regression must also be at this level. In fact, the model implies that this equation holds for any state (k^*, t), any d_t pair (d, d'), and any (appropriately chosen) renewal action d''. However, an estimation that covers all possible state and action combinations is not possible due to practical and computational limitations. To simplify this, for each state (k^*, t), I construct up to 85 possible action tuples (d, d', d'') and aggregate the data according to these choices. This is described in Appendix D.4.

Appendix D.4 also describes the data analogs for each of the components in Equation (14), e.g. school quality and distance measures. Calculating the payoffs requires making two assumptions:

Assumption 4: *The discount factor is assumed to be* $\beta = 0.9$.

Assumption 5: *The payoff to the outside option, exclusive of moving costs, is normalized to zero.*

$$\bar{u}(0, k_{it}, w_t) - MC(0, k_{it}) = 0, \forall k_{it}, \forall t$$

Assumption 4 is standard in the estimation of dynamic models as the discount rate cannot typically be identified.[26] Assumption 5 is also standard, as only differences in utility can be identified.

Next, I outline the estimation of Equation (14). The estimation is done separately by ethnic group, as the model parameterization allows all the coefficients to vary by type ψ.

5.4 School Related Variables

To deal with the endogeneity of school quality, I use the fact that the school components only enter the utility function for families with children who attend school. This suggests a natural strategy: I can compare two similar families who both have children, where one family's children are of school-age while the other family's children are not. In doing so, I can difference out all common factors – including unobserved neighborhood qualities – leaving only the school variables in the regression equation. A similar identification strategy was first used by Caetano (2019).

Consider Equation (14) for two families who have the same state variables *other* than the age/grade of their children. This means that they live in the same location, have the same tenure counter, and (as they are of the same type) have the same coefficient values. For a family without school-aged children, however, the school variables are all zero (Equation D.2). This means the difference in $Y_{d,d',d'',t}$ between a family with school-aged children and a family without school-aged children can be expressed as:

$$\begin{aligned}\Delta Y_{d,d',d'',t} &= Y_{d,d',d'',t}\{\text{School aged}\} - Y_{d,d',d'',t}\{\text{Non-school-aged}\} \\ &= (SQ_{dt} - SQ_{d't}) + (DS_d - DS_{d'}) - (SMC_{dt} - SMC_{d't}) \\ &\quad + (OOS_{dt} - OOS_{d't}) + \beta\left(SMC_{d,d'',t} - SMC_{d',d'',t}\right) \end{aligned} \tag{15}$$

To generate unbiased estimates on school valuation, the identifying assumption is that families with school-aged children value locations in the same way as families with non-school-aged children, except in terms of the school components. For example, families with a 7 year-old child (who is in primary school) should value neighborhood amenities such as parks and safety similarly to families with a 3 year-old child (who has not yet started school). A violation of the assumption would be if the neighborhood valuation changes discontinuously as children enter school or change school levels. I show evidence supporting this identification strategy in Appendix D.5.1.

[26] For reference, Diamond et al. (2017) set $\beta = 0.85$ and Davis et al. (2021) set $\beta = 0.95$ in similar residential choice models.

The estimation results are shown in Table D.2. These coefficients do not have a clear interpretation as they are in terms of utility relative to the outside option. They are also not comparable across ethnic groups. However, in section 6, I will convert them into economically meaningful WTP values. Reassuringly, the coefficients are estimated with the expected signs. School quality is generally viewed as a positive characteristic, while school moving cost is highly negative. The coefficient on distance varies in sign across the specifications. This is because DS represents the average distance between a home location and school, weighted by the probability of attending each school (as described in section 5.1). This means that the coefficient is only relevant for schools that a child has a *non-zero* probability of attending. Larger average distances could be viewed positively if they represent having access to a greater number of schools. Importantly, disutility towards distance is still being captured, albeit implicitly by the school attendance probabilities.[27]

5.5 Neighborhood Variables

Next, I estimate the coefficients on the observable neighborhood variables: price and ethnic shares.[28] For this, I use an instrumental variable (IV) approach to address the endogeneity.

To set up the regression, I use the estimated school coefficients from Equation (15) to find the residuals from Equation (14) after controlling for the school-related components. Let $\widetilde{Y}_{d,d',d'',t}$ represent these residuals. Then, using a similar strategy as Almagro and Domínguez-Iino (2021) and Scott (2013), I difference the expression for $\widetilde{Y}_{d,d',d'',t}$ across time to eliminate the time-invariant components. I consider two state-action tuples (k^*, t, d, d', d'') and (k^*, s, d, d', d''), where $t > s$. These tuples are similar on all dimensions except the calendar year. The equation to estimate is then:

$$\begin{aligned}\Delta_{ts}\widetilde{Y}_{d,d',d''} =& \widetilde{Y}_{d,d',d'',t} - \widetilde{Y}_{d,d',d'',s} \\ =& \alpha^r_\psi \left(\Delta_{ts} \ln r_d - \Delta_{ts} \ln r_{d'}\right) + \sum_{\varphi \in \Psi} \alpha^\varphi_\psi \left(\Delta_{ts} TS^\varphi_d - \Delta_{ts} TS^\varphi_{d'}\right) + \Delta_{ts}\epsilon_{dd'} \end{aligned} \tag{16}$$

where the operator Δ_{ts} represents a differencing across time between times t and s. The moving cost, tenure, and location fixed effects are differenced out because they are time-invariant in the utility function. This leaves the time-varying neighborhood variables: price and ethnic shares. The endogeneity problem persists however, as we could expect a correlation between $\Delta_{ts} \ln r$ or $\Delta_{ts} TS^\psi$ and the error term $\Delta_{ts}\epsilon$. Appendix D.5.2 shows the full derivation for Equation (16) and a further discussion on the endogeneity concern.

To instrument for price, I use a Hausman instrument commonly used in the industrial organization literature. I construct an index of non-neighboring prices for each location. The identifying assumption is that prices across the city are interrelated through market forces (relevance), but prices in non-

[27] As Figure D.5 shows, families are more likely to attend schools closer to their home.
[28] I create a house price index for each location and year (Appendix D.4). This uses both the sales and rental data.

neighboring areas should not have a direct relationship with local neighborhood amenities (satisfying the exogeneity restriction). This uses a similar intuition to the instrument of Bayer et al. (2007). An example and further details on the instrument construction are given in Appendix D.5.3.

To instrument for ethnic shares, I use the past settlement instrument. This is a shift-share instrument often used in the immigration literature. This relies on national-level shocks in the population of each ethnic group (the "shift") being independent from the group's neighborhood location distribution as of a baseline date (the "share"). I use data from the 1986 census as the baseline and rely on this identifying assumption to create a simulated ethnic share for each location and time. Appendix D.5.4 explains the instrument construction and discusses the instrument validity further.

The first stage results of the IV estimation are shown in Table D.3a. All the instruments are statistically significant and each regression has a large F-statistic. The second stage results are shown in Table D.3b. The price coefficient is always negative, as expected. I find that Polynesians are the most price sensitive, while whites are the least. This is consistent with the summary statistics of section 2.2. However, for the ethnicity shares, the point estimates are all negative, suggesting that all groups value areas with higher white shares. It may seem surprising that even non-white groups view their own ethnicity negatively. However, as described in Appendix A.2, these ethnic groups are broadly defined with heterogenous subgroups. These subgroups may not have positive preferences towards all those within their ethnic classifications (e.g. Chinese and Indian within Asian). Moreover, these coefficients are not precisely estimated, which could be a further reflection of heterogeneous preferences within each ethnic group.

5.6 Moving and Location Variables

The final step is to estimate the coefficients on the time-invariant moving and location variables. With the IV estimates, I can once again generate residuals after controlling for the price and ethnic share components. I then run a regression of the residuals on the remaining variables (location fixed effects, tenure, and moving costs). We should not expect these variables to be correlated with the unobserved errors, and so a standard OLS regression is sufficient. I show the results of this estimation in Table D.4. The results again show the expected sign. First, tenure is generally positively valued (however not for Asians, many of whom are recent immigrants and may not have as strong ties to the local communities). In contrast, moving costs are negative for all groups.

6 Discussion of Results

In this section, I provide the key takeaways of the structural model estimates. I focus on how parents respond to school quality and the school moving cost; the model estimates show that responses to both these factors vary across ethnicities and school levels. To find the full effect that schools have on residential segregation, I simulate the market under counterfactual utility specifications.

6.1 School Quality Valuation

First, I discuss the valuation of school quality. I use the price elasticity estimates (Table D.3b) to convert each ethnicity's marginal utility of school quality (Table D.2) into a willingness to pay (WTP) figure.[29] The interpretation of this value is the percentage increase in annual house price that a family is willing to pay for an increase of one decile score per child per year. The average annual house price in 2019 is NZ$28,387 (US$20,723).[30]

I also benchmark the structural estimates to the reduced form estimates from section 3.1. As discussed in section 3.1.2, these results are not directly comparable. Namely, the reduced form estimates give the WTP for a one decile score increase without accounting for the time horizon that this additional school quality will be enjoyed over. In Appendix E.1, I describe how I normalize the reduce form results. This relies on using the family panel to generate statistics about families who move into the boundary areas. I argue that the key statistic is the present discounted value of school years that a family expects to experience when they move to a new home. Using this figure, I can convert the BDD sales results (Table 1a) into a per-child-per-year valuation.

Given these calculations, I plot the WTP estimates in Figure 6. I find that white families are willing to pay 0.5% higher annual house prices for an increase in one primary school decile score, per child per year. In contrast, their valuation of high school is more than double with a WTP of 1.1% for high school quality.[31] As with the reduced form estimates, I find that their preference for intermediate schools is statistically indistinguishable from zero. The pattern for Polynesian and Asian families is remarkably similar, with the highest WTP for intermediate school (1.5% and 1.4%, respectively), followed by primary school (1.1% for both), and finally high school (0.7% and 0.6%, respectively). However, the estimates do not have enough precision to confidently conclude whether this reflects their true prioritization of school levels.

The most striking feature of Figure 6 is that the reduced form estimates are most closely aligned with the

[29] For a coefficient estimate $\widehat{\theta}^V_\psi$ and price elasticity estimate $\widehat{\alpha}^r_\psi$, the WTP is calculated as $\left(\exp\left(-\widehat{\theta}^V_\psi/\widehat{\alpha}^r_\psi\right)-1\right)\times 100$.

[30] This is calculated using the house price index for each of the model's locations, weighted by their population.

[31] For a one standard deviation increase in quality (i.e. 3 decile scores), this corresponds to an average additional housing cost of NZ$433 and NZ$937 per child per year for primary and high school, respectively. As a reference point, annual tuition fees for private schools start at approximately NZ$25,000.

structural estimates for white families. This could be driven by the fact that the reduced form estimates are based on house sales, while white households are disproportionately more likely to be home-owners (Table A.5). If these results hold true in other settings, then it calls into question the generalizability of the boundary discontinuity design. Even if the BDD estimates were accurate, this exercise shows that their interpretation is limited without further information about the family composition and housing tenure.

Figure 6: Willingness to Pay for School Quality

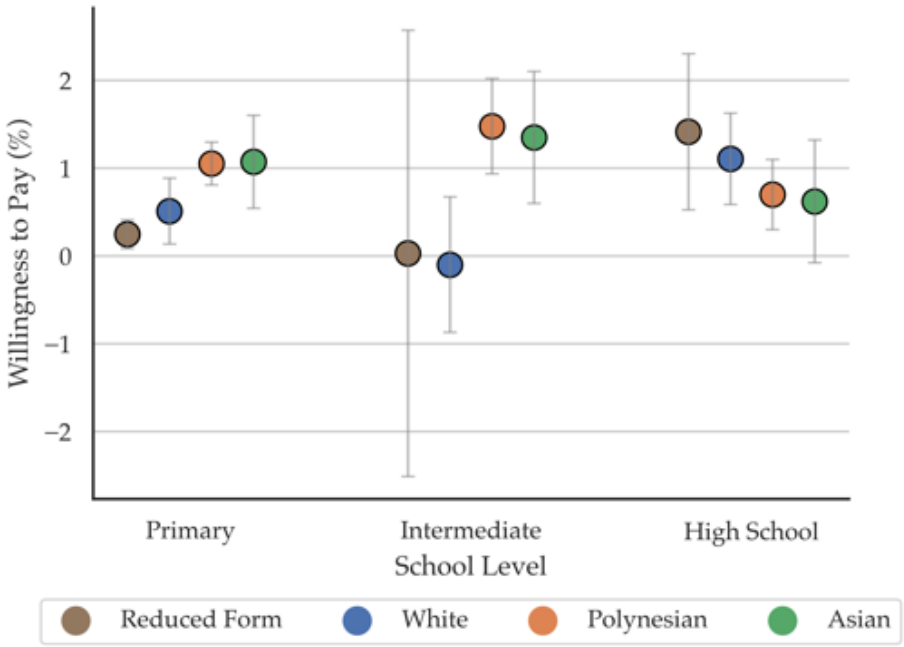

Note: The figure shows the willingness to pay estimates for an increase of one decile score, per child per year. The reduced form estimates are from a boundary discontinuity design (Table 1a, column 2), scaled using an estimated present discounted value of school years for the given school level. The ethnic group results are from the structural model estimation (Table D.2), scaled by the price sensitivity estimates (Table D.3b). 95% confidence intervals are plotted for each estimate.

6.2 School Moving Cost

Following the steps in section 6.1, I convert the estimates on school moving costs (Table D.2) into WTP values. The results are shown in Figure 7, where I plot the estimates of the school moving cost by ethnicity and school level. As these values are all negative, we can interpret this as the change in house price that is required to offset the disutility of changing one child's school. There are two key takeaways from this graph.

The first takeaway is that all ethnicities have an aversion to changing their child's school. White families, in particular, exhibit the greatest disutility with a WTP ranging from –14.3% for primary school to –9.2% for high school. The WTP of Polynesian families, in contrast, is almost half that of white families in each

school level, e.g. their estimated WTP is –6.8% and –4.5% for primary and high school, respectively. Asian families, despite being the most mobile group, have a WTP that is in between the other two groups. The Asian families exhibit a WTP of –10.3% for primary school and –6.7% for high school. The second takeaway is that the aversion to changing schools decreases as the child is older. This pattern holds true for all ethnicities. For each group, the WTP for primary school is approximately 1.5 times that of the high school WTP.

Figure 7: School Moving Cost Estimates

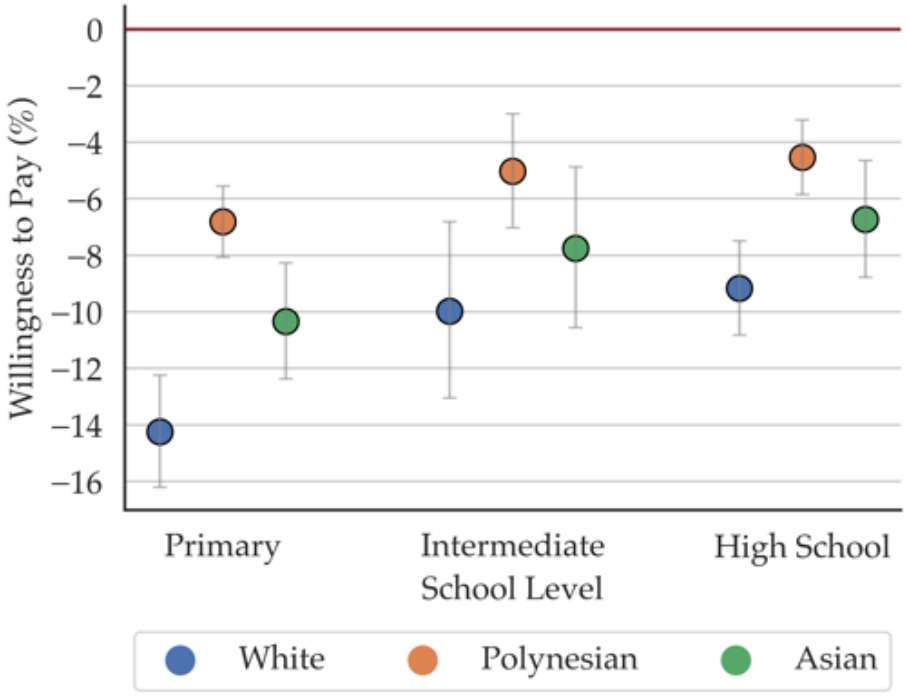

Note: The figure shows the estimated school moving cost. Each point represents the willingness to pay for changing one child's school within the same school level. The estimates are from the structural model estimation (Table D.2), scaled by the price sensitivity estimates (Table D.3b). 95% confidence intervals are plotted for each estimate.

The existing literature has found high moving costs for housing (Bayer et al., 2016; Diamond et al., 2017; Almagro and Domínguez-Iino, 2021), and this is an importance source of the dynamics in the residential choice problem.[32] I show the WTP estimates for the other components of utility in Figure E.1. We can see that the fixed costs of moving are almost equal to the school moving costs for each ethnic group. This means that schools may be explaining half of the moving costs for families. Past studies have not focused on identifying the source of the moving costs. However, understanding why moving costs are large is important for predicting the efficacy of policies. For example, we may see low take-up of housing vouchers if – as in this case – parents are unwilling to move houses due to the cost of interrupting their children's education. Policy makers could then instead target vouchers to families whose children have

[32]The results in Almagro and Domínguez-Iino (2021), whose model is most similar to mine, imply that the fixed cost of moving has a WTP of 30-40% in annual budget (income net of rent) for most of their distinct groups, with some groups having moving costs of almost 90% WTP. This does not include the per distance moving costs.

not yet started school to increase take-up and welfare.

The school moving cost finding is also relevant for the education literature, where researchers seek to better understand what parents value about schools (Abdulkadiroglu et al., 2020). These results show that quality is only part of the story. To give a sense that these school moving costs are substantial, we can ask whether parents can offset these costs by simply changing to a higher quality school. In Appendix E.2, I show through a simple example that a white family with a first grader would have to move to a school that is more than 2 standard deviations higher in quality to offset their school moving cost. These significant frictions have policy implications. If parents are reluctant to change their child's school, this lowers the competitive forces under a school choice policy. As a result, these costs may inhibit school choice policies from raising school quality.[33]

6.3 Counterfactuals

To understand the impact schools have on residential segregation, I simulate the market using the estimated parameters. In general, counterfactuals based on the ECCP estimator require strong assumptions (Kalouptsidi et al., 2020). In my setting, simulating the full dynamic model has a high computational burden. To simplify, I set $\beta = 0$ to convert it into a static problem. In Appendix E.3, I provide further details on this simplification and the calibrations I make to account for the dynamic factors.

I consider two counterfactual utility specifications. The first counterfactual sets the coefficient on school quality to be equal to zero (*No Quality*). This can be thought of as having equal school quality in all neighborhoods. The second counterfactual sets the coefficients on all school components to zero (*No School*). This counterfactual represents a world in which schools no longer matter for a family's residential choice. Relative to the first counterfactual, the second counterfactual highlights the additional effect of school moving costs.[34]

To evaluate changes in residential segregation, I calculate the dissimilarity index (DI) in each period. Logan and Stults (2021) classify changes in the DI (in units of 0.01 increments) over a decade as follows: differences of less than 5 units indicate small to no change in segregation; 5-10 units indicates moderate change; and above 10 units indicates a very significant change in residential segregation. I also calculate the DI from the benchmark simulation (no changes to utility) and from the data (using the family panel).

The results are shown in Figure 8. Panel (a) shows the Polynesian-white DI. The benchmark model – despite being a simplification of the full dynamic model – closely tracks the data's DI. In the benchmark, the Polynesian-white DI increases by 3.6 units over the sample period. The DI increases even more in the counterfactuals: by 6.4 units in the first, and by 10.2 units in the second. This suggests that

[33] Implementing school choice will not necessarily reduce the school moving costs – the New Zealand system already offers a high degree of school choice and yet these costs remain high.

[34] Under this counterfactual, there is also no preference over distance to school. I find that the effect of school distance in the simulations to be negligible, because the school attendance probabilities already capture the households' disutility to distance.

as schools become less relevant to the family's residential choice, segregation between Polynesian and white families increases. For the Asian-white DI, I find a similar pattern of increasing segregation, but at a smaller magnitude. In the benchmark, the DI changes by –1.4 units (decreasing segregation). In the counterfactuals, the DI is unchanged in the first and increases by 1.9 units in the second.

Figure 8: Dissimilarity Index in Counterfactual Simulations

(a) Polynesian-White

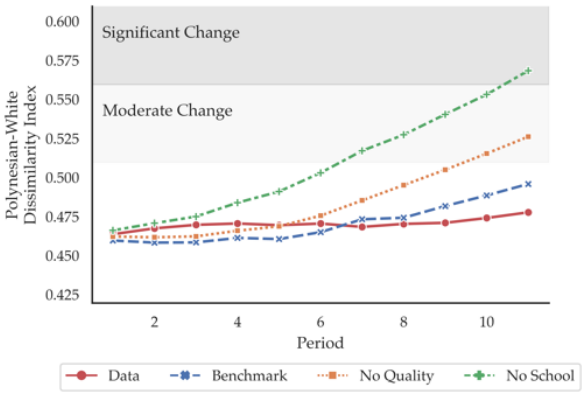

(b) Asian-White

Note: The figures show the dissimilarity index in each period (representing the sample period of 2008-2019). The index is calculated using the Level 3 area units (Figure A.16c).

Figure 9: Change in Home Location's White Share Among White Movers

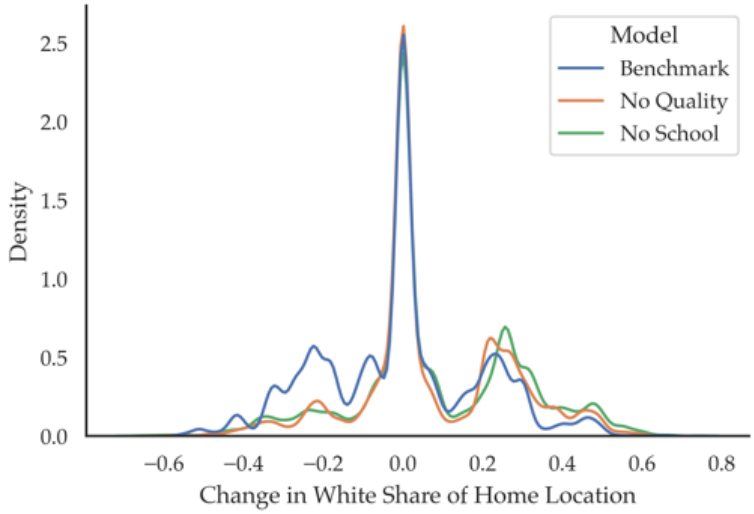

Note: The figure shows a kernel density estimation for each simulation. For white families who move within the city, the horizontal axis shows the difference in the white share between their previous home location and their new home location. The distribution is taken over all moves observed in each counterfactual simulation.

To understand what is driving this change in segregation, I check whether families are moving into areas with a higher share of their own ethnicity. In Figure 9, I focus on within-city moves by white families. The distribution shows the difference in white share between their previous home location and their

new home location. In the counterfactuals, this distribution has a greater weight on positive differences. In other words, white families are more likely to move to whiter areas. Caetano and Maheshri (2019) also find that moving costs reduce residential segregation: if families could move more, they would segregate further. Here, both school quality and school moving costs prevent white families from moving more often. In contrast, I do not find these effects for Polynesian (Figure E.2) and Asian (Figure E.3) families. These results are likely driven by the estimated preferences over ethnic shares (Table D.3b), and so the results may differ in other contexts.

7 Conclusion

This paper studies how schools affect where families choose to live in a city and, as a result, whether schools have an impact on residential segregation. I develop a dynamic model of residential choice, which incorporates children and schools into the family's problem. Using administrative microdata on families in Auckland, New Zealand, I construct a novel panel dataset that links parents to children and tracks their home location and the children's schools over a 12 year period. With this data, I estimate the model parameters separately by ethnic group to assess heterogeneity in the effects of schools.

The estimation generates three main findings. First, school quality matters to families, however the valuation varies by ethnicity and school level. I leverage the existence of catchment areas to compare the model's estimates against ones derived from a boundary discontinuity design. I find that the reduced form estimates are aligned only with white families' valuations of quality. More generally, without microdata on the family's structure and moving behavior, reduced form estimates do not provide a standardized measure of valuation. Second, I find that families experience high costs to moving homes if that results in their child changing schools. These costs are larger for younger children and are unlikely to be offset by changing to a higher quality school. Third, through a simplified counterfactual, I show that schools reduce residential segregation. Without schools, families would move more often, with white families choosing to segregate further.

Children play a key part in a family's choice of neighborhoods. Importantly, their effect on moving decisions acts in a dynamic way – families change with the birth of children and children's schooling needs change as they grow up. This paper shows that children and schools should be incorporated into dynamic models of residential choice. The findings also highlight that preferences are heterogeneous across ethnicities. Researchers should account for this – or seek out settings where this variation can be identified – to strengthen the generalizability of their results. Moreover, capturing this heterogeneity is critical for understanding the drivers of residential segregation. While this paper has focused on studying segregation along ethnic lines, another critical (and closely tied) issue is that of income segregation. Future work should incorporate the parents' employment into the model to then estimate how work impacts the home location choice. With this, researchers can then study the interplay between income and ethnic residential segregation, while also accounting for the role of children and schools.

References

Abdulkadiroglu, Atila, Parag A Pathak, Jonathan Schellenberg, and Christopher R Walters, "Do parents value school effectiveness?," *American Economic Review*, 2020, *110* (5), 1502–1539.

Almagro, Milena and Tomás Domínguez-Iino, "Location Sorting and Endogenous Amenities: Evidence from Amsterdam," 2021.

Arcidiacono, Peter and Robert A Miller, "Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity," *Econometrica*, 2011, *79* (6), 1823–1867.

Auckland Regional Council, "A brief history of Auckland's urban form," Technical Report 2010.

Avery, Christopher and Parag A Pathak, "The distributional consequences of public school choice," *American Economic Review*, 2021, *111* (1), 129–152.

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

_ , Robert McMillan, Alvin Murphy, and Christopher Timmins, "A dynamic model of demand for houses and neighborhoods," *Econometrica*, 2016, *84* (3), 893–942.

Bergman, Peter, Raj Chetty, Stefanie DeLuca, Nathaniel Hendren, Lawrence F Katz, and Christopher Palmer, "Creating moves to opportunity: Experimental evidence on barriers to neighborhood choice," Technical Report 2019.

Black, Sandra E and Stephen Machin, "Housing valuations of school performance," in "Handbook of the Economics of Education," Vol. 3, Elsevier, 2011, pp. 485–519.

Boterman, Willem R, Sako Musterd, and Dorien Manting, "Multiple dimensions of residential segregation. The case of the metropolitan area of Amsterdam," *Urban Geography*, 2021, *42* (4), 481–506.

Caetano, Gregorio, "Neighborhood sorting and the value of public school quality," *Journal of Urban Economics*, 2019, *114*, 103193.

_ and Vikram Maheshri, "A Unified Empirical Framework to Study Segregation," Technical Report 2019.

Davis, Morris A, Jesse Gregory, and Daniel A Hartley, "The long run effects of low income housing on neighborhood composition," Technical Report 2018.

_ , Jesse M Gregory, Daniel A Hartley, and Kegon T K Tan, "Neighborhood effects and housing vouchers," Technical Report 2021.

Diamond, Rebecca, Tim McQuade, and Franklin Qian, "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality," in "Proceedings. Annual Conference on Taxation and Minutes of the Annual Meeting of the National Tax Association," Vol. 110 JSTOR 2017, pp. 1–72.

Epple, Dennis and Richard E Romano, "Competition between private and public schools, vouchers, and peer-group effects," *American Economic Review*, 1998, pp. 33–62.

Friesen, Wardlow, "The demographic transformation of inner city Auckland," *New Zealand Population Review*, 2009, *35* (1), 55–74.

Grbic, Douglas, Hiromi Ishizawa, and Charles Crothers, "Ethnic residential segregation in New Zealand, 1991-2006," *Social science research*, 2010, *39* (1), 25–38.

Ho, Elsie, "The changing face of Asian peoples in New Zealand," *New Zealand Population Review*, 2015, *41* (95), 95–118.

Hotz, V Joseph and Robert A Miller, "Conditional choice probabilities and the estimation of dynamic models," *The Review of Economic Studies*, 1993, *60* (3), 497–529.

Jones, Kelvyn, Ron Johnston, James Forrest, Chris Charlton, and David Manley, "Ethnic and class residential segregation: exploring their intersection - a multilevel analysis of ancestry and occupational class in Sydney," *Urban Studies*, 2018, *55* (6), 1163–1184.

Kalouptsidi, Myrto, Paul T Scott, and Eduardo Souza-Rodrigues, "Linear IV regression estimators for structural dynamic discrete choice models," *Journal of Econometrics*, 2020.

Kramer, Michael R, "Residential segregation and health," *Neighborhoods and health*, 2018, pp. 321–356.

Logan, John R. and Brian Stults, "The Persistence of Segregation in the Metropolis: New Findings from the 2020 Census," Technical Report, Diversity and Disparities Project, Brown University 2021.

Logan, Trevon D and John M Parman, "The national rise in residential segregation," *The Journal of Economic History*, 2017, *77* (1), 127–170.

Massey, Douglas and Nancy A Denton, *American apartheid: Segregation and the making of the underclass*, Harvard university press, 1993.

Monarrez, Tomas E., "School Attendance Boundaries and the Segregation of Public Schools in the US," *Working Paper*, 2021.

Nechyba, Thomas, "School finance, spatial income segregation, and the nature of communities," *Journal of Urban Economics*, 2003, *54* (1), 61–88.

OECD, "Balancing School Choice and Equity: An International Perspective Based on Pisa," Technical Report oct 2019.

PPTA, "New Zealand Schools: The Decile System," Technical Report 2013.

Raval, Devesh, Ted Rosenbaum, and Steven A Tenn, "A semiparametric discrete choice model: An application to hospital mergers," *Economic Inquiry*, 2017, *55* (4), 1919–1944.

Rust, John, "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher," *Econometrica*, 1987, pp. 999–1033.

Scott, Paul T, "Dynamic Discrete Choice Estimation of Agricultural Land Use," 2013.

Tomorrow's Schools Independent Taskforce, "Our Schooling Futures: Stronger Together," Technical Report 2018.

Zheng, Angela, "The Valuation of Local School Quality under School Choice," 2021.

Appendices

A Setting Appendix

A.1 Details on the New Zealand Educational System

My analysis focuses on public schools in Auckland. Approximately, 84% of students in Auckland attend public schools.[35] In 2019, there were 353 public schools in Auckland. 272 of these schools can be categorized as a primary (grades 1-6), intermediate (grades 7-8), or high school (grades 9-13). There are also combined primary and intermediate schools (grades 1-8; 72 schools in 2019), combined intermediate and high schools (grades 7-13; 6 schools), as well as schools with all grade levels (grades 1-13; 3 schools). For my analysis, I treat these combined schools as separate for each of their respective school levels.

A.1.1 Key Features of the New Zealand System

As discussed in section 2.1, there are three key features that distinguish the New Zealand public educational system. These features make this an ideal setting to study the research questions.

The first feature is that it is a decentralized system. There are no school districts; instead, schools are autonomous entities that are governed by an elected board of trustees. Public schools receive funding from the national government and supplement this through donations. In contrast to the United States, schools are not financed through local property taxes. This is an advantage of the setting as having school funding tied to property taxes could complicate any analysis related to residential sorting.

The second feature is that it is a system of complete choice. Students can choose to apply to any public school regardless of where they live. However, acceptance is not always guaranteed. To prevent oversubscription, schools can establish a catchment area, which is commonly referred to as a school zone. This is a geographic boundary that determines which residences receive guaranteed admission into the school.[36] Any student who lives inside a school's zone (in-zone student) must receive admission into the school if they choose to apply there. Any student who lives outside of the zone (out-of-zone student) can still apply to the school, but they are given lower priority in the admissions process.[37] Only schools that are oversubscribed – or are close to being oversubscribed – can introduce a school zone. In contrast, schools without a zone ("unzoned schools") must accept all students who wish to attend.

[35] 6% of students attend private schools, and 10% attend "special character" public schools (these are former private schools that have been integrated into the public system; most of these are Catholic schools). Both of these school types charge attendance fees and have different admissions processes than the standard public schools.

[36] Figures A.2 and A.3 show school zones in Auckland for some primary and high schools, respectively. Note that school zones tend to be smallest for primary schools and largest for high schools.

[37] After admitting all in-zone applicants, the school calculates the number of seats available for out-of-zone students. If there are more out-of-zone applications than seats, the school holds a stratified lottery to determine admission.

Figure A.1: Fraction of Schools Zoned

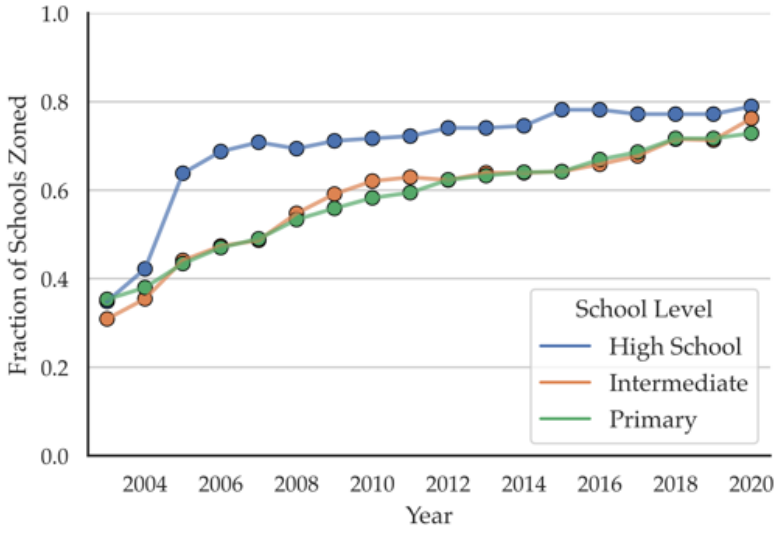

Note: The figure shows the fraction of public schools in Auckland that have a school zone. Fractions are calculated within each school level, in each year. Data is from the New Zealand Ministry of Education.

Over time, more schools have become "zoned", i.e. introduced a school zone. Figure A.1 shows the fraction of schools at each school level. Empirically, once a school becomes zoned, it does not ever become unzoned. However, the fraction of zoned schools may fall over time due to schools opening or closing. As of 2019, the number of zoned schools is: 185 of 258 (72%) of primary schools, 87 of 122 (71%) intermediate schools, and 44 of 57 (77%) of high schools. The presence of school zones allows me to study the valuation of school quality using a reduced form framework (section 3), which I can then use to benchmark the structural estimates (section 6). Moreover, this admissions process means that the link between housing and schools is much weaker in New Zealand than it is in in other countries. If schools are tied to residential segregation in this setting, then it is plausible that the connection is even stronger in areas where residential choice and school admissions are even more intertwined.

The final feature is the measure of school quality. In New Zealand, there are no standardized tests or test scores that parents can use to compare schools. Instead, a commonly used measure of quality is a government index called the decile score. The national government assigns every school in the country a decile score of 1 through 10, which reflects the socioeconomic status of the student body. A lower decile score indicates that the school draws its students from higher poverty areas. The decile score is *not* a measure of school effectiveness (e.g. value-added); rather, the government uses this score when allocating its funding to schools. However, the general public commonly uses (or misinterprets) the decile score as a school quality measure.[38] A key advantage of the decile scores is that they are inputs-

[38] A government taskforce stated in their 2018 report: "School deciles were introduced in 1995 as a funding mechanism, but decile is now often viewed as a proxy for the quality of teaching and learning" (Tomorrow's Schools Independent Taskforce, 2018). Similarly, a 2013 report by the secondary school teachers' union stated that "school decile has become such a factor in the (mis)interpretation by parents of the quality or status of a school" (PPTA, 2013).

based measure of school quality, as it is entirely a function of which students attend a school. This allows us to abstract away from typical considerations of what determines school quality, such as the quality of teachers and classroom sizes. Another advantage is that they are calculated using a relative ranking of all schools in the country, regardless of their school level. This means that the decile score is directly comparable across all school levels.[39] This will allow for a more meaningful comparison when assessing whether school quality valuation changes by school level. Appendix A.1.2 below gives further details on how decile scores are calculated.

Figure A.2: School Zones in Auckland (Primary Schools)

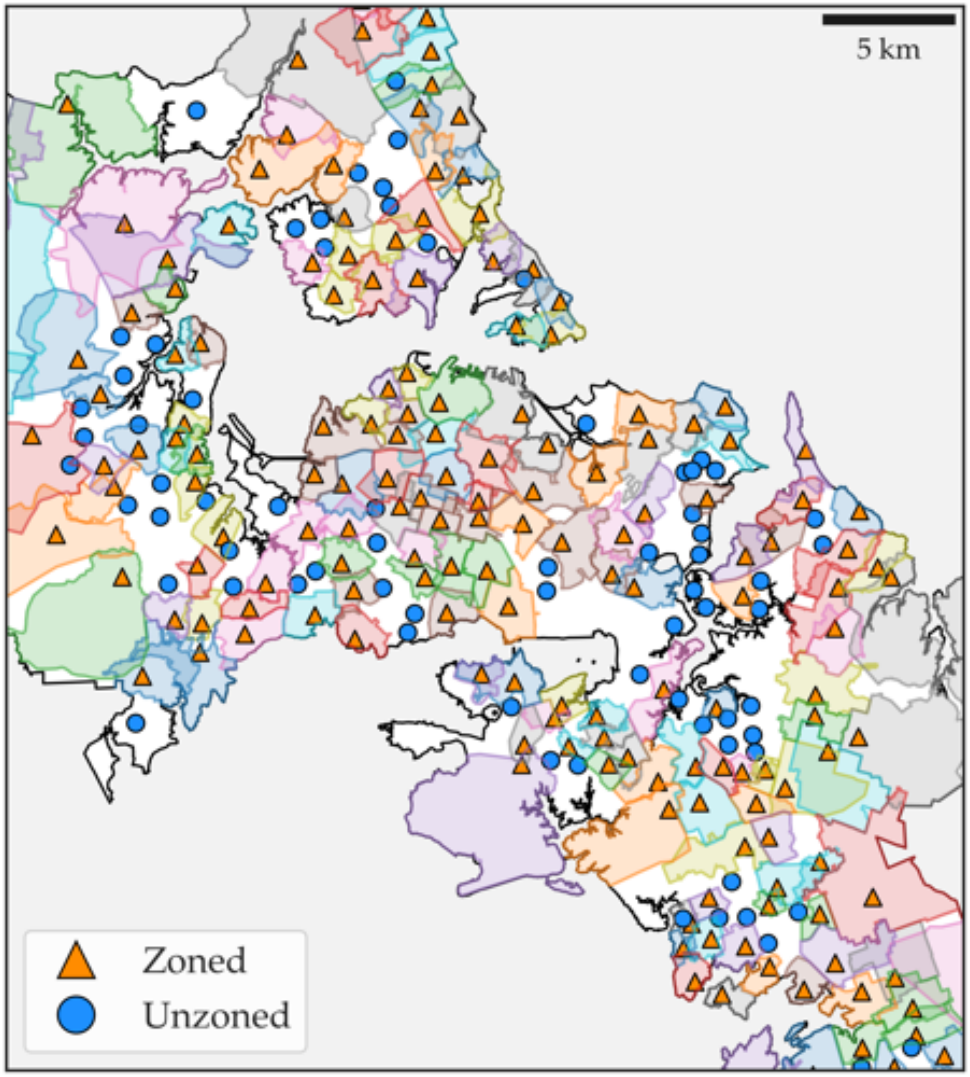

Notes: The maps show the location of primary schools for a section of the Auckland urban area. Schools that are zoned are represented by the orange triangles. Schools that are unzoned are represented by the blue circles. The figure overlays the school zones (as of 2018), which have been assigned random colors.

[39] An alternative measure, such as test scores, may not have that advantage. For example, test scores for an elementary school may not be comparable to test scores for a high school, even after a within-level normalization.

Figure A.3: School Zones in Auckland (High Schools)

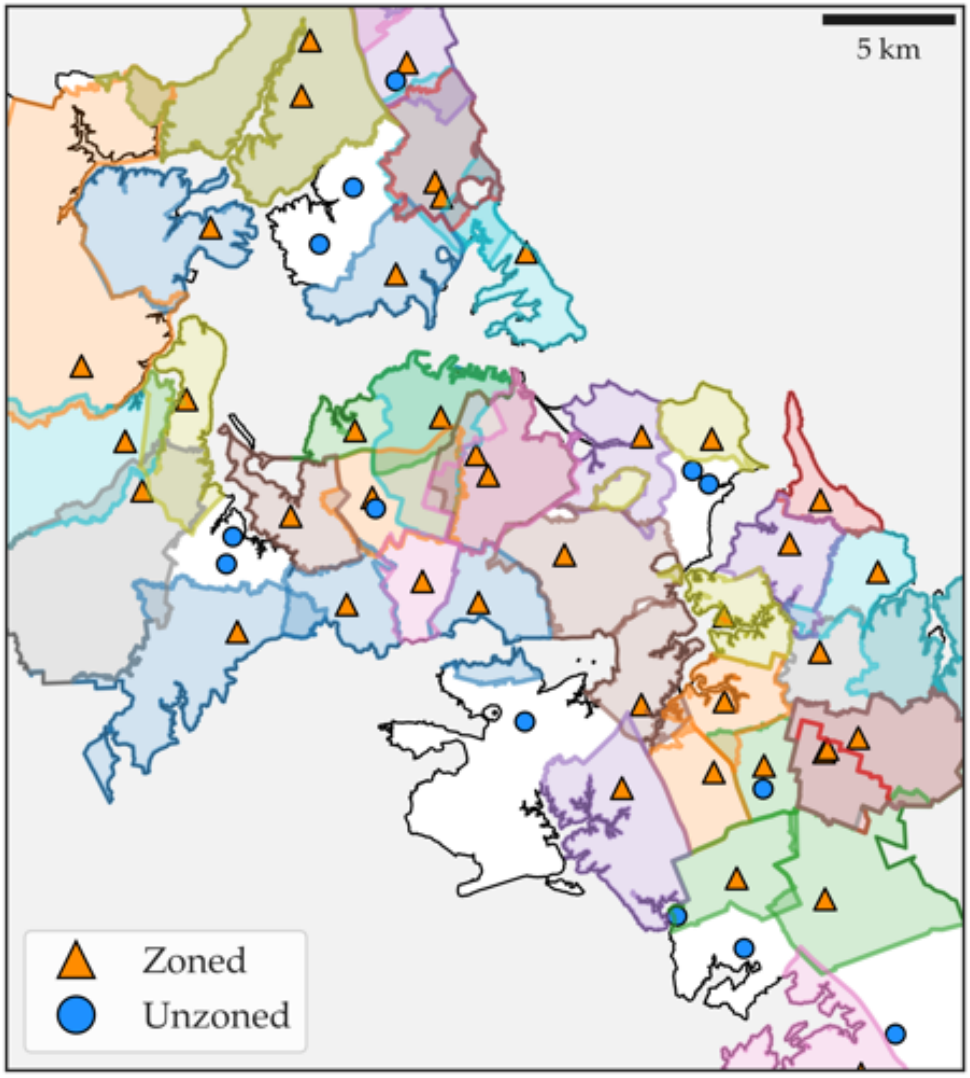

Notes: The maps show the location of high schools for a section of the Auckland urban area. Schools that are zoned are represented by the orange triangles. Schools that are unzoned are represented by the blue circles. The figure overlays the school zones (as of 2018), which have been assigned random colors.

A.1.2 Decile Scores

Each public school in the country is designated a decile score (an integer from 1 through 10). The schools with the lowest decile score (1) have students drawn from the lowest socioeconomic status (SES) areas in the country and the highest decile (10) schools have students drawn from highest SES areas. To calculate this score, the national government uses the following process. Each census meshblock m (analogous to a U.S. census block) is given a SES score ρ_m based on 5 equally weighted factors:

1. percentage of parents in meshblock m who are employed in a low-skill job (p_1)
2. percentage of parents in meshblock m who are receiving income support (p_2)
3. percentage of parents in meshblock m without formal qualifications (p_3)
4. percentage of households in meshblock m with an income in the bottom quintile (p_4)
5. percentage of households in meshblock m which experience crowdedness (p_5)

where $\rho_m = \sum_{i=1}^{5} p_i$ ranges from 0 to 500, with higher scores indicating lower SES. For each school j, let the fraction of students living in meshblock m be s_{mj}. Each school j is assigned an SES score $\bar{\rho}_j$ by weighting each meshblock's SES score by the fraction of students who live there:

$$\bar{\rho}_j = \sum_{m=1}^{M} s_{mj}\rho_m$$

The decile score is then the decile of the school's SES score $\bar{\rho}_j$ within the national distribution. Figure A.4 shows this process graphically. In (a), the raw score $\bar{\rho}_j$ is plotted for the national distribution in 2015. In (b), the assigned decile score is shown relative to the school's raw score, where a higher $\bar{\rho}_j$ indicates a lower decile score. The vertical lines indicate the decile cutoffs, which are determined ex post and change every time the deciles are re-calculated. The decile scores are generally updated every five years after the national census, however, each year schools may request a review of their score in if they believe their student composition has changed dramatically.[40]

Figure A.5 shows the distribution of decile scores for schools in Auckland. By definition, the national distribution has an approximately equal number of schools within each decile band. In contrast, Auckland schools tend to be skewed as having both very low and very high decile scores. This shows that this is setting with large heterogeneity and already indicates the presence of segregation (in this case, along socioeconomic lines).

[40]The most recent censuses have been conducted in 2006, 2013, and 2018. A census was planned for 2011 but was delayed by two years due to the 2011 Christchurch earthquake.

Figure A.4: Decile Score Calculation (2015)

(a) Raw Score Distribution

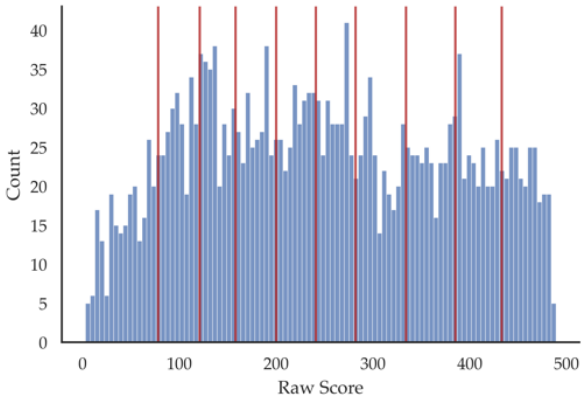

(b) Decile Score Assignment

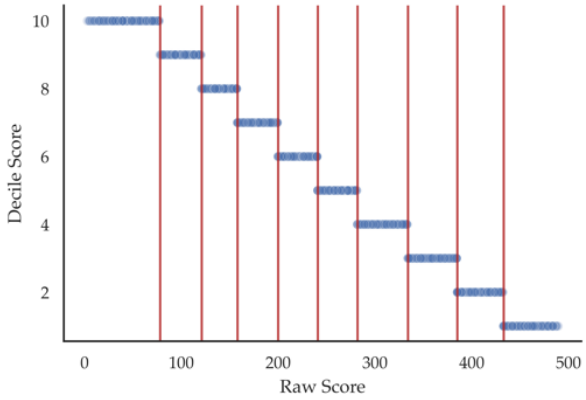

Note: The figures show how decile scores are calculated. Each school in the country is assigned a score out of 500 (sum of five percentages), where 0 is indicates highest SES and 500 indicates lowest SES. The distribution of this score is shown in panel (a). Cut-offs are determined to group schools into approximately 10 even groups (i.e. deciles). Schools with the lowest raw measure receive a decile score 10, while schools with the highest raw measure receive a decile score of 1. All schools in the same score band receive the same decile score. This is shown in panel (b).

Figure A.5: Distribution of Decile Scores in Auckland

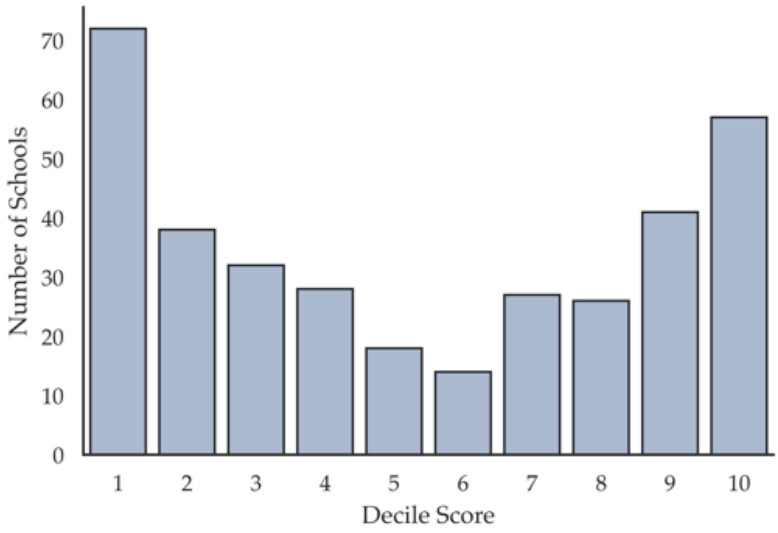

Note: The figure shows the number of Auckland schools in each decile score as of 2015. All school levels (primary, intermediate, high school) are included.

A.2 Defining Ethnic Groups

In my analysis, I designate three distinct ethnic groups: white, Polynesian, and Asian. However, as ethnicity is a social construct, this classification first merits some discussion.

The New Zealand government currently defines five major ethnic groups, where ethnicity refers to a person's cultural affiliation and not race, ancestry, or citizenship. The groups are: European, Māori, Pacific Islanders, Asian, MELAA (Middle Eastern, Latin American, or African), and Other. European includes New Zealand Europeans (the majority group in the country), as well as direct migrants from Europe. Māori are the indigenous people of New Zealand. Pacific Islanders (or Pacific Peoples) refers to those ties to the neighboring island nations in Polynesia, such as Samoa, Tonga, and Fiji. Asian includes all of Asia except the Middle East, with the largest groups identifying as Chinese, Indian, and Filipino. Importantly, people can self-identify with multiple ethnicities and a person's ethnic identification may change over time in government records.[41]

To be consistent with the model (where a family's type does not change over time), I create an ethnic grouping that is distinct and stable. I define "white" as people who identify as European but not as Māori, Pacific Islander, or Asian. "Polynesian" are those who identify as either Māori or Pacific Islander, but not Asian. "Asian" are those who identify as Asian (and possibly other ethnicities as well). All other people who do not fit into these groups are excluded from the remainder of the analysis. While I use these labels as a shorthand for ethnicity, in reality, these groups are not completely distinct and still exhibit substantial within-group heterogeneity.

I apply my ethnicity categorization to microdata from the 2018 New Zealand census. In Table A.1, I show the demographic breakdown of individuals living in Auckland. Under my categorization, whites are the largest group in the city (42.2% of the population), followed by Asians (30.6%), and then Polynesians (24.8%).[42] The table also shows how the (distinct) ethnic groups that I created compare to the government's (overlapping) ethnic categories. Most notably, 32% of Polynesian people identify as European (i.e. white). This overlap is primarily driven by people who identify as both Māori and New Zealand European.[43] Similarly, Table A.4 shows that the majority of Asian parents were born in Asia (77.5%); however, a substantial minority (13.2%) were born in the Pacific Islands. This is explained by Indo-Fijians, who tend to identify as both Pacific Islanders (Fijian) and Asian (Indian). Under my classification, these families are categorized as Asian, though this further highlights the heterogeneity within each group.

Finally, it is also important to note that the statistics in Table A.1 reflect the city as of 2018. Auckland's

[41] For the panel construction, I use the Stats NZ personal demographic data. This contains information on ethnicity identification that takes into account that reported identities can change. An individual is recorded as identifying with an ethnic group if they have ever identified with that group at any point in time across any of the datasets collected by Stats NZ.

[42] 2.4% of the city does not fall into these groups. This represents the population excluded from my sample.

[43] 53% of those who identify as Māori also identify as European. For Pacific Islanders, 17% also identify as European.

ethnic composition has also changed over time (my sample period is 2008 to 2019). In particular, the Asian population has been growing each year – in large part due to immigration – with the Asian share of the city being 19.8% in the 2006 census, and 24.8% in the 2013 census.

Table A.1: Official Ethnic Identification by Ethnic Groups

	Ethnic Groups			
	White	Polynesian	Asian	Other
Ethnic Identification in 2018 Census				
European	100.0	31.9	6.2	0.0
Māori	0.0	45.3	1.5	0.0
Pacific Islander	0.0	63.9	2.9	0.0
Asian	0.0	0.0	100.0	0.0
Middle Eastern/Latin American/African	0.8	0.4	0.4	80.7
Other	1.8	0.5	0.8	19.5
Number of Individuals	629,586	370,272	456,150	36,558
Share of City (% of total population)	42.2	24.8	30.6	2.4

Note: The table shows how individuals living in Auckland per the 2018 census identified their ethnicity according to the official Statistics New Zealand categories (rows in first panel). The columns represent the ethnic groups I create for my analysis. Each cell in the first panel indicates the percentage of individuals in my ethnic classification that identify with the official ethnic category. The ethnic groups I create (columns) are distinct groups by construction, while people can identify with multiple of the official ethnic groups (rows). This means columns in the first panel range from 0% to 100%, but can sum to over 100%. In accordance with Stats NZ confidentiality rules, the number of observations has been randomly rounded to a multiple of 3 (RR3).

A.3 Measuring Residential Segregation

There are a number of possible indices that researchers can use to measure residential segregation, each with advantages and drawbacks. For each segregation measure, the city is divided into L neighborhoods indexed by $\ell = \{1, ..., L\}$. Consider two population sub-groups a and b. In a location ℓ, the population of each group is a_ℓ and b_ℓ, respectively. Let n_ℓ be the total population in location ℓ and $N = \sum_\ell n_\ell$ be the total population of the city. Finally, let the total population of each group in the city be denoted as $A = \sum_\ell a_\ell$ and $B = \sum_\ell b_\ell$.

The most widely used approach in the literature is the dissimilarity index. For two groups a and b, the dissimilarity index (D_{ab}) is calculated as:

$$D_{ab} = \frac{1}{2} \sum_{\ell=1}^{L} \left| \frac{a_\ell}{A} - \frac{b_\ell}{B} \right| \tag{A.1}$$

The interpretation of D_{ab} is the proportion of group a that would need to move in order to have a spatial distribution even to that of group b. This means numbers closer to 1 indicate greater segregation, and numbers closer to 0 indicate lower segregation. This value is symmetric, i.e. $D_{ab} = D_{ba}$, and must be calculated using exactly two groups. This means it is limited to measuring how segregated group a is *relative* to group b (and vice versa).

Another common approach is the isolation index, which captures the probability of interaction between the groups. Using the same notation as above, the isolation index (I_{ab}) is calculated as:

$$I_{ab} = \sum_{\ell=1}^{L} \left(\frac{a_\ell}{A} \right) \left(\frac{b_\ell}{n_\ell} \right) \tag{A.2}$$

I_{ab} is interpreted as the probability that a member of group a will interact with a member of group b. Here, numbers closer to 1 indicate a higher probability of interaction and therefore lower segregation between the two groups. As with the dissimilarity index, the isolation index can only compare two groups at a time (however, it is generally not symmetric).

As this setting has three major ethnic groups, the final measure of segregation I use is the entropy index. For a location i, let $p_{\ell\psi}$ be the share of type ψ in location ℓ (where $\psi \in \{a, b, c\}$). The entropy index h_ℓ for location ℓ is calculated as:

$$h_\ell = -\sum_{\psi=1}^{\Psi} p_{\ell\psi} \ln(p_{\ell\psi}) \tag{A.3}$$

This is interpreted as a measure of concentration in each location ℓ.[44] In this case, as there are three groups, h_ℓ ranges from 0 (only one group in the location) to $\ln 3 \approx 1.10$ (equal shares of each group). In Figure A.6, I plot the entropy measures h_ℓ for locations ℓ in the city.

To calculate a city wide measure of entropy, we first calculate $\hat{H}$, a city-wide measure of entropy (using Equation A.3, where $\ell = city$).[45] Next, we calculate $\bar{H}$, which is the population weighted average of h_ℓ:

$$\bar{H} = \sum_{\ell=1}^{L} \frac{n_\ell}{N} h_\ell$$

The entropy index (H) is then calculated as:

$$H = \frac{\hat{H} - \bar{H}}{\hat{H}} \tag{A.4}$$

The entropy index is calculated by comparing the observed level of entropy $(\bar{H})$ to the level of entropy if the city was perfectly integrated $(\hat{H})$, i.e. if each location had the same composition as the entire city.

[44] For cases with locations with zero shares of a subgroup, i.e. $p_{\ell\psi} = 0$, let $p_{\ell\psi} \ln(p_{\ell\psi}) = 0$

[45] Using the shares shown in Table A.1 (excluding other), $\hat{H} = 1.07$ for Auckland in 2018.

Figure A.6: Local Entropy Measures

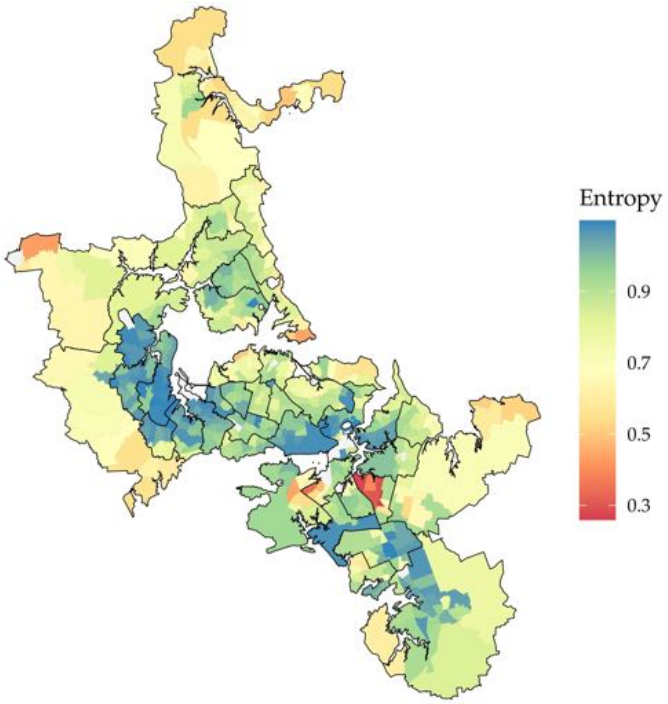

Note: The map shows in each area (Level 2 units, Figure A.16b) the entropy measure of the ethnic shares. Values closer to $ln(3) \approx 1.1$ indicate equal distribution of all three ethnic groups. I overlay the borders of the locations used in the structural model estimation (Figure A.16c).

This means that $H = 0$ indicates perfect integration and $H = 1$ indicates perfect segregation.

An important point to consider across all of these cases is that these measures depend on the size of the locations $\ell \in \{1, ..., L\}$. To show a range of possible values, I re-calculate each of these indices based on the three smallest geographic units (Levels 1 to 3) shown in Figure A.16. These indices in the U.S. are often calculated using census tracts, which are comparable to the Level 2 geographic unit. The structural model estimation (section 5) uses Level 3 as the geographic unit, which generally produces lower measures of residential segregation. This suggests that any index may not fully capture the intra-neighborhood levels of segregation. Table A.2 shows the different segregation measures using each of the three indices as well as the different geographic units.

Interpreting the segregation indices can often be difficult. A useful benchmark is to compare against measures found in other cities. Using the Diversity and Disparities database, I find the segregation indices of U.S. cities as of 2020.[46] San Diego, California is a comparable city to Auckland. It has a similar population of 1.3 million, with a racial/ethnic breakdown of 41% white, 30% Hispanic, and 21% Asian. In San Diego, the dissimilarity index is 0.521 for Hispanic-white (comparable to Auckland's Polynesian-white index) and 0.415 for Asian-white (comparable to Auckland's Asian-white index). To continue

[46]This database categorizes the population of each city into the distinct groups of: non-Hispanic white; non-Hispanic Black; non-Hispanic Asian; Hispanic; and other. I omit the "non-Hispanic" qualifier when describing statistics for the U.S. cities.

the comparison started in section 2.2, the dissimilarity index in New York City is 0.477 for Hispanic-white and 0.42 for Asian-white. For Los Angeles, these figures are 0.607 and 0.378, respectively. The demographic shares in New York are 31% white, 28% Hispanic, and 17% Asian. For Los Angeles, these are 29%, 47%, and 13%, respectively. While their demographic compositions are different, these are also close to the Auckland measures of segregation.

Table A.2: Measures of Residential Segregation

	Dissimilarity Index		**Isolation Index**		**Entropy**
Geo. Level	Polynesian-White	Asian-White	Polynesian-White	Asian-White	All
1	0.586	0.458	0.260	0.344	0.261
2	0.545	0.388	0.280	0.374	0.200
3	0.493	0.305	0.304	0.400	0.146

Note: The table shows different specifications for measuring residential segregation. Each column represents a segregation index. The dissimilarity and isolation index are always comparisons of two groups. I show the Polynesian-White and Asian-White index values. The entropy index is a single city-wide measure that captures segregation across all groups. Each row represents a different geographic unit level, as shown in Figure A.16. I re-calculate each index for each geographic levels 1 to 3.

A.4 Panel Data Construction Process

Stats NZ's Integrated Data Infrastructure (IDI) is a database that hosts a wide array of administrative datasets.[47] Importantly, Stats NZ assigns each individual a unique identifier across all the datasets within the IDI. This extensive – and interconnected – database allows me to construct the panel required for the structural estimation.

My focus is on identifying the sample population relevant to estimating the model. This means identifying families with children who have ever lived in Auckland. I first start by finding all children who live in Auckland (via the address history dataset) or attend an Auckland school (via the school records) at any point during my sample period of 2008 to 2019. Using the birth records and the censuses, I then identify their parents. Here, "parent" can refer to a biological parent (via the birth records) or the adult that they reside with and have indicated as their parent (via the censuses). This process may identify multiple parents, so I use the address history file to find the parents that the child is most often living with during the sample period. With this, I categorize families as sets of parents and children who live together. I only keep families for whom I can identify the parents and who have no more than four children.[48]

I use the mother's home location as the family's location. This is because my goal is on knowing where

[47]The full list of data in the IDI is available at the Statistics New Zealand website.

[48]The restriction on number of children excludes less than 1% of families, according to Table A.4, but simplifies the data management and analysis.

Figure A.7: Living with Parents

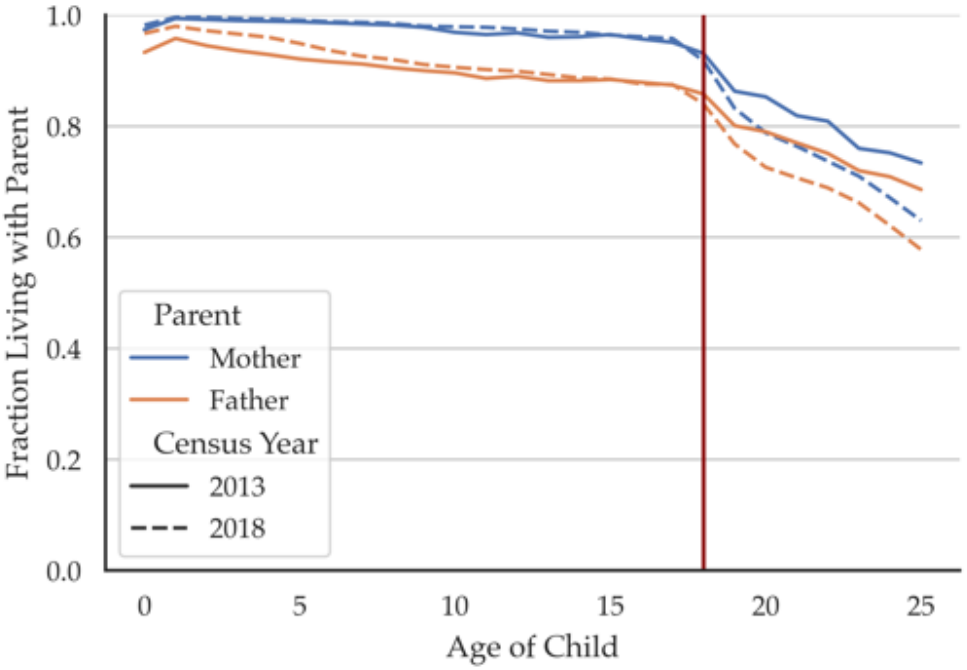

Note: The figure shows the proportion of children in my sample data who are living with their identified mother and identified father as of the two most recent censuses (2013 and 2018). I calculate this proportion separately by census year and the age of the child at the time of the census. Home locations for each child and parent is per their response in the census. These statistics are only among families who are able to be matched to the census data.

the children live. I find that children are much more likely to be living with their identified mother than with their identified father (Figure A.7).[49] I further restrict so that families also must have been living in Auckland for at least one period in the sample period. For each family, I am able to identify their home location for every period in the sample period 2008 to 2019, including before any children are born.[50] I assign each family one location per year, using the location that has the longest residence spell in a given year. For confidentiality reasons, the data in the IDI does not show the exact home address. It is listed at the meshblock level, a geographic unit that is analogous to a census block in the U.S.[51]

Finally, I assign each family an ethnicity. This is not straightforward as families may have different ethnicities: not only can parents have multiethnic identities, but inter-ethnic relationships can result in children identifying with more ethnicities than each of their parents do. I assign each family as identifying with one of the five official ethnic groups (European, Māori, Pacific Islanders, Asian, MELAA, Other) if either the mother identifies with the group or if more than half of the children identify with

[49] In addition, I am almost always able to identify a child's mother, while children's fathers are not always identified or children within the same family can have different fathers.

[50] While my analysis is restricted to Auckland, the address history dataset allows me identify all locations that the family is living in, even outside of Auckland. Using information on border arrivals and departures in the IDI, I can also identify when families enter and leave New Zealand. All those living outside New Zealand in a given period are identified as "overseas". This is important due to high levels of immigration during my sample period, especially among Asian families.

[51] Each meshblock contains approximately 30-50 households and is represented as the Level 1 unit in Figure A.16a.

the group. I then assign each family one ethnic group (white, Polynesian, Asian) using the classification approach described in section A.2.

Table A.6 shows that the my family panel contains 302,508 unique families and 3,288,381 family-year observations. While other studies focus on cities larger than Auckland, the scope of their data is typically more limited (which means that their data also covers a smaller percentage of their city's population). For example: Davis et al. (2021) have 1,787,558 person-year observations over a 15 year period; Bayer et al. (2016) have 220,403 observations over an 11 year period; Almagro and Domínguez-Iino (2021) have 614,410 individuals over an 11 year period, though the panel is unbalanced and they do not report the number of households (e.g. 54.5% of their individuals are children); Caetano (2019) has 153,102 families in a cross-sectional dataset. Studying Auckland also gives me a large enough sample size on each of the major ethnic groups, which ensures diversity in my data (see footnote 3).

In Figure A.8, I show the total number of families by ethnicity in my sample dataset. As expected, I find that the number of Asian families is growing rapidly each year, such that it is equal to the number of white families in the final period.[52] Table A.6 provides summary statistics about the families in the sample; these statistics align with the census statistics in Tables A.4 and A.5. This shows that the sample construction is accurately reflecting the population of interest.

Figure A.8: Annual Count of Auckland Families in Data

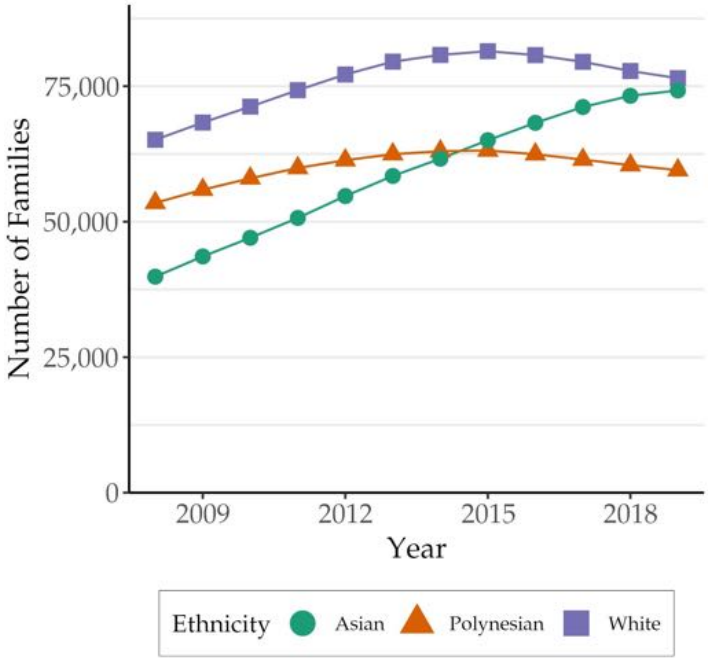

Note: This figure shows the number of families (by ethnic group) in the sample data that are living in Auckland each year. Families are included in the sample from 3 years prior to the birth of their eldest child up to when their youngest child turns 25. The sample is a panel dataset I construct using administrative records from Stats NZ.

[52]While white families make up the largest number of people in the city (Table A.1), they are also the least likely to be in families (Table A.3).

Table A.3 shows that 57.2% of the Auckland population is part of a family, either as a parent (27.6%) or as a child (24.5%). While my analysis focuses on the residential choice of families, this shows that these choices determine the home location for a majority of the city's residents.

Table A.3: Family Role of Population (2018)

	Ethnic Group			
	White	Polynesian	Asian	**All**
Role in Family (%)				
Parent	27.6	24.0	30.4	27.6
Child	24.5	38.2	29.9	29.7
Not in Family	47.9	37.8	39.7	42.8
Total Proportion	100.0	100.0	100.0	100.0
Number of Individuals	629,586	370,272	456,150	1,456,008

Note: The table shows what role individuals living in Auckland per the 2018 census play in a family. A family is defined as a household of parents and children. Each cell in the first panel shows the proportion of people in the ethnic group who are either in a family as a parent, as a child, or not in a family. These proportions sum to 100%. In accordance with Stats NZ confidentiality rules, the number of observations has been randomly rounded to a multiple of 3 (RR3).

A.5 Data Validity Checks

I discuss three validity checks for the data construction process.

First, I check whether the sample has sufficient coverage of Auckland families. Figure A.9 compares the number of children living in Auckland in my sample data to the official number of children reported to be living in Auckland per the census. This comparison is within each child's age group and for each of the three most recent censuses.[53] In general, coverage is relatively high with over 75% coverage for most years and age groups. However, the older cohorts in the earlier years have worse coverage, dropping to almost 50% in the 16-17 year-old age groups. This is for two reasons. First, identifying the children in the older cohorts is harder because the school records appear to be of lower quality in the earlier years.[54] Second, identifying their parents is also more challenging. This is because: (i) the proportion of people linked to their birth records decreases in older cohorts, and (ii) many of these children are no longer living with their parents by the time of the 2013 census. As stated above, I exclude children from my

[53] My sample period begins in 2008 and there is a census in 2006. In each period, I need to know the family's previous home for the model analysis. This means that I also identify each in-sample family's home location in 2007. Therefore, I compare the child counts in my sample as of 2007 to the counts of the 2006 census as this is the closest comparison point.

[54] Quality here refers to the extent to which individuals in the schools records can be matched to other datasets (e.g. the census). This relies on Statistics New Zealand assigning the correct identifier using the reported personally identifiable information (PII), which may be limited by the quality of the PII. The total number of students in each school-year aligns with publicly available enrollment counts.

Figure A.9: Children Counts in Data versus Census Counts

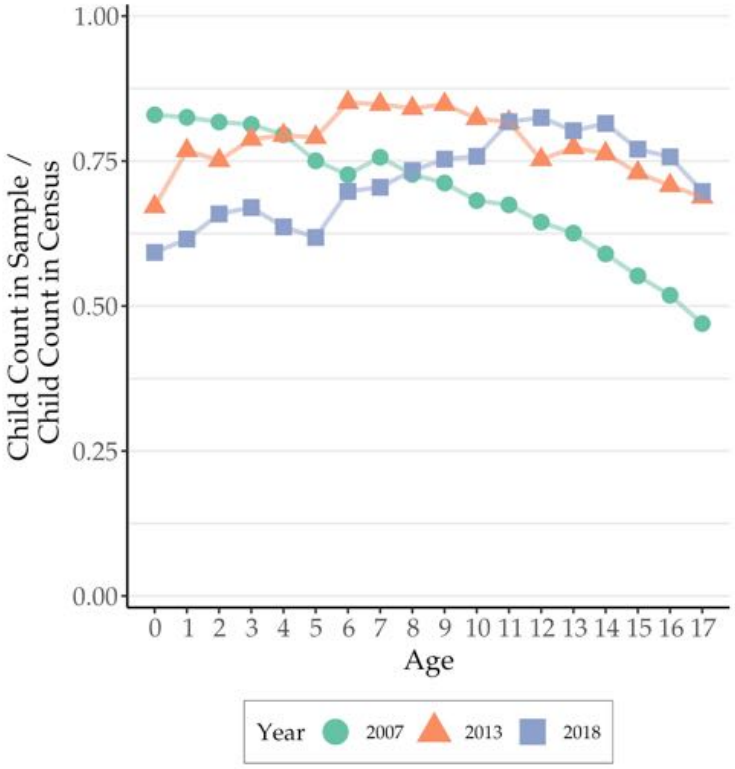

Note: The figure shows the number of children identified in my sample divided by the number of children in the official census counts (vertical axis), by age group (horizontal axis). This is restricted to families who are living in Auckland. The three most recent census were in 2006, 2013, and 2018. My period of analysis is from 2008 to 2019, though my sample data also includes 2007 information too. I compare the 2013 and 2018 sample statistics to their respective censuses. I compare the 2007 sample statistics to the 2006 census.

sample if their parents cannot be identified (as discussed in section 2.3). This explains why I observe lower coverage for these cohorts. Fortunately, it appears that the coverage has improved significantly by 2013. Though I cannot conduct this exercise for years between 2007 and 2013, it is likely that the lower coverage is only an issue for a few cohorts in the initial years of my sample period.

Second, I check whether the identified home location is likely to be accurate. Figure A.10 compares the distance between the child's school and their home location. Note that the home location is sourced from the identified mother's address history records, while the school is from the child's school records. These independent datasets are being connected using the parent-child links I have identified. We see that children in the sample are far more likely to be attending schools that are closer to them, which is to be expected. Moreover, this pattern is stronger for primary schools than for high schools, which is also consistent with expectations. This suggests that the sample construction process is likely to be capturing the family's home location quite well. Moreover, it also suggests that the parent-child links I have identified are also correct.

Figure A.10: Distribution of Distances to School

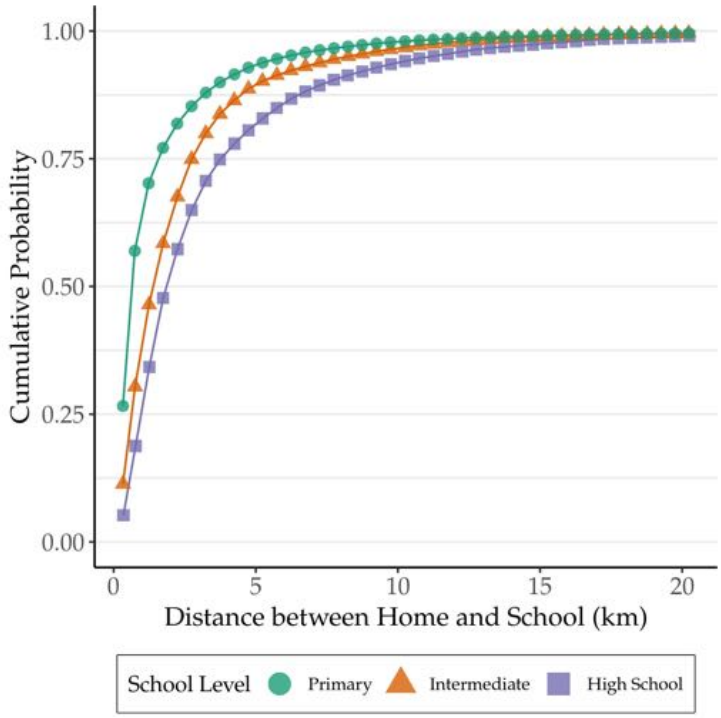

Note: The figure shows the cumulative distribution of distances between a family's identified home location (Level 1 unit, Figure A.16a) and the location of a child's school. Distances are straight-line distances in kilometers (km) between the Level 1 centroid and the school coordinates. The distribution is shown separately for each school level (primary, intermediate, and high school). Each unit is a child-year.

Finally, there may be a concern that the ethnicity classifications may lead to inaccuracies. One potential issue is that I assign each family into one ethnic group, but this may not reflect the actual ethnic proportions we see in terms of individuals. Figure A.11 compares the ethnic shares in each location using families as the unit of observation against the ethnic shares when using individuals as the unit of observation. The figure shows that these two measures are highly correlated. Therefore, studying the shares at the family-level is likely to be a good proxy for understanding segregation more broadly at the individual-level.[55] Another issue relates to the multi-ethnic identities that my classification does not capture. The most notable concern is that there are a high number of individuals who identify as both European and Māori. I classify these people as Polynesian, which could raise concerns if they are better classified as white. To show that this is not the case, I look at the residential distribution of individuals who identify as both white (European) and Polynesian (Māori or Pacific Islander) using the 2018 census. Figure A.12 shows that these individuals are more likely to be found in areas with a high proportion of people who identify *only* as Polynesian rather than in areas with a high proportion of people who identify *only* as white. Therefore, categorizing white-Polynesian families with the other Polynesian families seems to be the appropriate classification.

[55] An argument can also be made that family-level segregation is the more appropriate measure. Individual-level measures treat children equally to adults, but children typically do not choose their home location independently from their parents.

Figure A.11: Ethnicity Shares at Individual versus Family Level

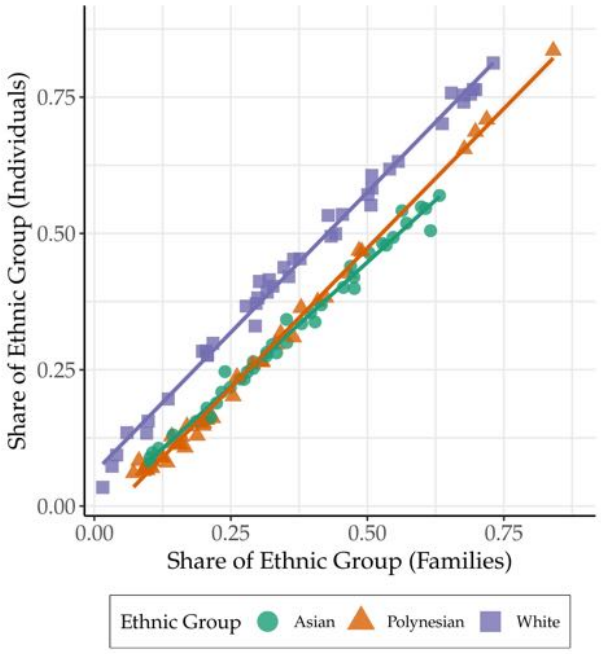

Note: The figure shows the ethnic share statistic in each location when calculated at the individual-level using census data (vertical axis) versus at the family-level using my sample panel data (horizontal axis). All statistics are as of 2018 and locations are set as the Level 3 geographic unit (Figure A.16c).

Figure A.12: Correlations with White-Polynesian Shares

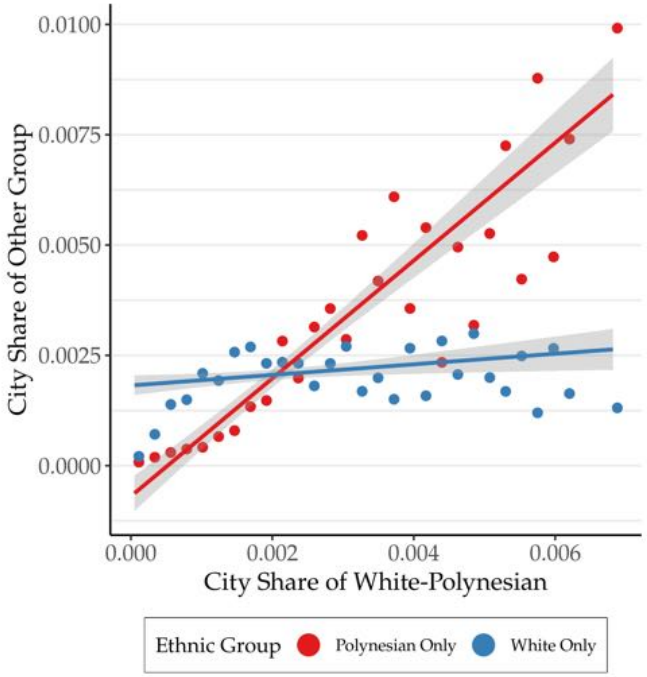

Note: The figure shows where individuals who identify as both Polynesian (Māori or Pacific Islander) and white (European) are most likely to live. On the horizontal axis, I plot the share of Polynesian-white individuals living in each location. On the vertical axis, I plot the share of individuals who identify only as Polynesian (red) and only as white (blue) living in that same location. Shares are calculated over each group's population in the city. I plot a binned scatterplot and linear regression (with 95% confidence interval). Statistics are calculated using the 2018 census and locations are set as the Level 2 area unit (Figure A.16b).

A.6 Supporting Figures

Figure A.13: Auckland Region Map

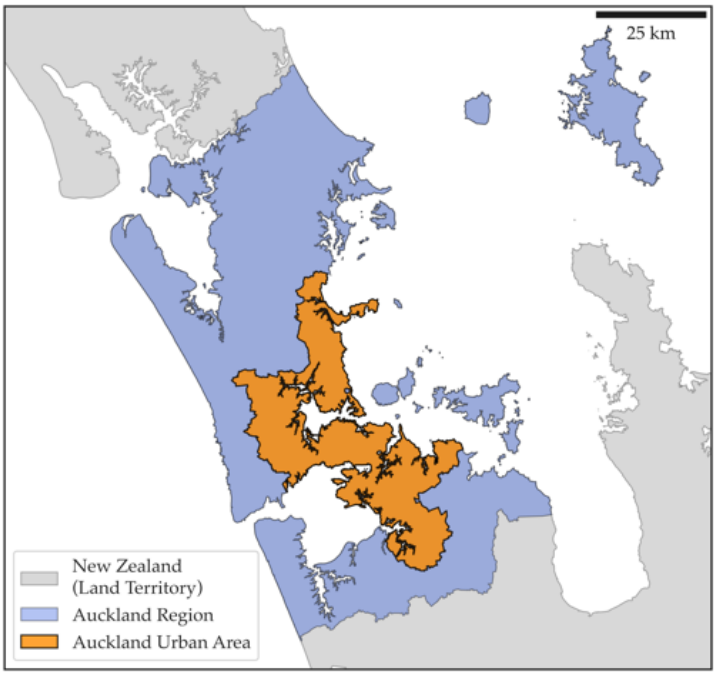

Note: The map shows the location of the Auckland region (blue) within New Zealand (gray). The urban area on the major island (orange) is the geographic area of analysis.

Figure A.14: Population Distribution in Auckland (2018)

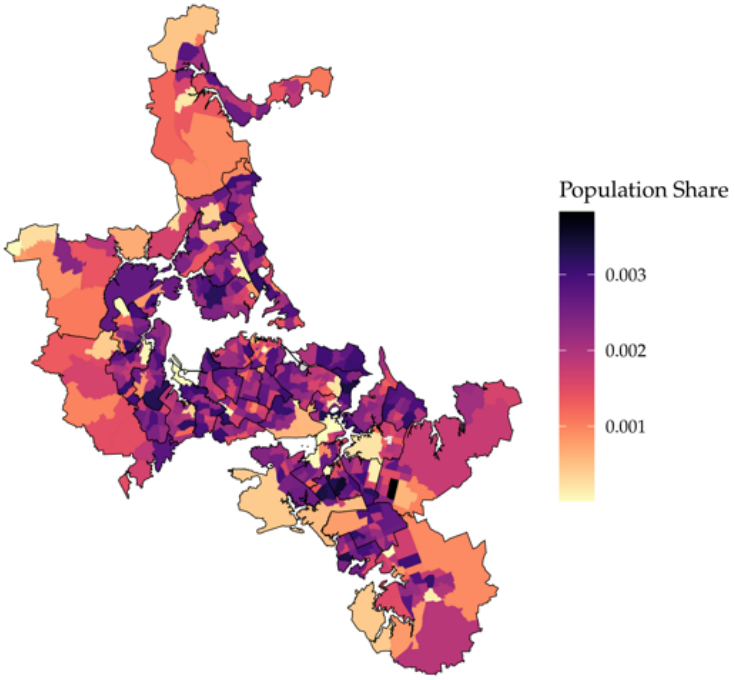

Note: The figure shows the population distribution within the Auckland urban area using the 2018 census. I shade each location (Level 2 unit, Figure A.16b) according to the fraction of the city's total population that is living in that location.

Figure A.15: House Prices versus Ethnic Shares

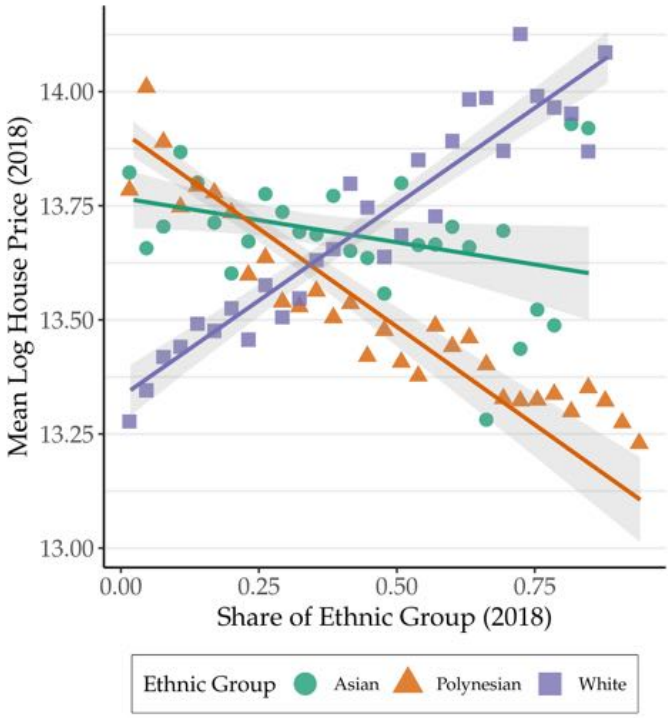

Note: The figure shows how average house prices (in log New Zealand dollars) vary by share of each ethnic group. Statistics are as of 2018 using the 2018 census to calculate ethnic shares. Locations are the Level 2 area unit (Figure A.16b). The figure shows a binned scatter plot with a line of best fit (and 95% confidence interval).

A.7 Supporting Tables

- Table A.4 shows statistics about the parents of families in Auckland using the 2018 census microdata.
- Table A.5 shows statistics about families in Auckland using the 2018 census microdata.
- Table A.6 shows statistics about families in the family panel data I construct.

Table A.4: Summary Statistics of Parents

	Ethnic Group		
	White	Polynesian	Asian
Demographics			
Age (mean)	47.3	43.9	44.0
	(0.027)	(0.044)	(0.030)
Employed†	0.817	0.657	0.698
	(0.001)	(0.002)	(0.001)
Female†	0.533	0.585	0.548
	(0.001)	(0.002)	(0.001)
Has university degree†	0.422	0.143	0.465
	(0.001)	(0.001)	(0.001)
Immigration Status			
Born in New Zealand†	0.660	0.573	0.044
	(0.001)	(0.002)	(0.001)
Born in Pacific Islands†	0.001	0.380	0.132
	(0.000)	(0.002)	(0.001)
Born in Asia†	0.005	0.001	0.775
	(0.000)	(0.000)	(0.001)
Immigrated within last 5 years†	0.051	0.026	0.180
	(0.001)	(0.001)	(0.001)
Years in New Zealand (mean)	20.3	25.4	14.0
	(0.068)	(0.076)	(0.025)
Years in New Zealand (median)	15	24	14
	–	–	–
Housing Tenure			
Years in current home (mean)	8.7	9.0	5.7
	(0.023)	(0.038)	(0.018)
Years in current home (median)	6	5	4
	–	–	–
Same home location as one year ago†	0.837	0.708	0.755
	(0.001)	(0.002)	(0.001)
Same home location as five years ago†	0.507	0.415	0.349
	(0.001)	(0.002)	(0.001)
Number of Individuals (parents)	173,637	88,716	138,795

Note: The table shows summary statistics for Auckland adults who are parents of dependent children. Statistics are calculated using responses to the 2018 census. † indicates a mean of an indicator variable (i.e. interpreted as a proportion). Standard errors are in parenthesis below each estimated mean. In accordance with Stats NZ confidentiality rules, the number of observations has been randomly rounded to a multiple of 3 (RR3).

Table A.5: Summary Statistics of Families

	Ethnic Group		
	White	Polynesian	Asian
Family Composition			
Number of children (mean)	1.8	2.3	1.7
	(0.003)	(0.006)	(0.003)
1 child†	0.434	0.360	0.482
	(0.002)	(0.002)	(0.002)
2 children†	0.412	0.293	0.396
	(0.002)	(0.002)	(0.002)
3 children†	0.125	0.175	0.095
	(0.001)	(0.002)	(0.001)
4 children†	0.023	0.093	0.020
	(0.000)	(0.001)	(0.001)
5 or more children†	0.006	0.079	0.007
	(0.000)	(0.001)	(0.000)
Single Parent Household†	0.239	0.432	0.183
	(0.001)	(0.002)	(0.001)
Household Income			
Income below NZ$50,000†	0.113	0.385	0.260
	(0.001)	(0.002)	(0.002)
Income between NZ$50,000-$100,000†	0.215	0.300	0.330
	(0.001)	(0.002)	(0.002)
Income above NZ$100,000†	0.673	0.315	0.410
	(0.002)	(0.002)	(0.002)
Current Residence			
Number of bedrooms (mean)	3.5	3.4	3.5
	(0.003)	(0.005)	(0.004)
Has mortgage†	0.517	0.251	0.475
	(0.002)	(0.002)	(0.002)
Renter†	0.283	0.675	0.349
	(0.001)	(0.002)	(0.002)
Weekly rent in NZ$ (median)	550	400	490
	–	–	–
Number of Families	98,205	55,926	78,876

Note: The table shows summary statistics for Auckland adults who are parents of dependent children. Statistics are calculated using responses to the 2018 census. † indicates a mean of an indicator variable. Standard errors are in parenthesis below each estimated mean. Dollar amounts are in New Zealand dollars (1 NZD = 0.73 USD). Income indicates self-reported total household income. Weekly rent calculated among renters. In accordance with Stats NZ confidentiality rules, the number of observations has been randomly rounded to a multiple of 3 (RR3).

Table A.6: Summary Statistics of Sample Families

	Ethnic Group		
	White	Polynesian	Asian
Age and Tenure			
Age in 2018 (mean)	45.0	42.9	41.5
	(0.024)	(0.030)	(0.028)
Immigrated during sample period†	0.082	0.033	0.298
	(0.001)	(0.001)	(0.002)
Years in current home as of 2018 (mean)	8.1	8.3	5.4
	(0.022)	(0.027)	(0.020)
House tenure as of 2018 (median)	6	6	3
	–	–	–
Family Composition			
Number of children: Mean	1.8	2.0	1.7
	(0.002)	(0.003)	(0.002)
1 child†	0.390	0.404	0.479
	(0.001)	(0.002)	(0.002)
2 children†	0.441	0.330	0.409
	(0.001)	(0.002)	(0.002)
3 children†	0.144	0.177	0.095
	(0.001)	(0.001)	(0.001)
4 or more children†	0.025	0.089	0.018
	(0.000)	(0.001)	(0.000)
Age of youngest child in 2018 (mean)	12.2	12.2	10.4
	(0.022)	(0.025)	(0.027)
House Moves			
Number of moves in sample period (mean)	1.4	1.6	1.9
	(0.004)	(0.006)	(0.005)
0 moves†	0.393	0.373	0.198
	(0.001)	(0.002)	(0.001)
1 move†	0.232	0.205	0.263
	(0.001)	(0.001)	(0.001)
2 moves†	0.170	0.154	0.233
	(0.001)	(0.001)	(0.001)
3 moves†	0.103	0.110	0.158
	(0.001)	(0.001)	(0.001)
4 or more moves†	0.102	0.157	0.148
	(0.001)	(0.001)	(0.001)
Observation Counts			
Number of families	123,951	89,316	89,241
Number of family-year observations	1,372,260	1,024,974	891,147
Number of observations with house moves	169,350	145,554	168,927

Note: The table shows summary statistics for families in the sample data, which is a panel of families who ever lived in Auckland between 2008 and 2019. † indicates a mean of an indicator variable. Standard errors are in parenthesis below each estimated mean. In accordance with Stats NZ confidentiality rules, the number of observations has been randomly rounded to a multiple of 3 (RR3).

A.8 Geographic Units

Figure A.16: Geographic Units

(a) Level 1

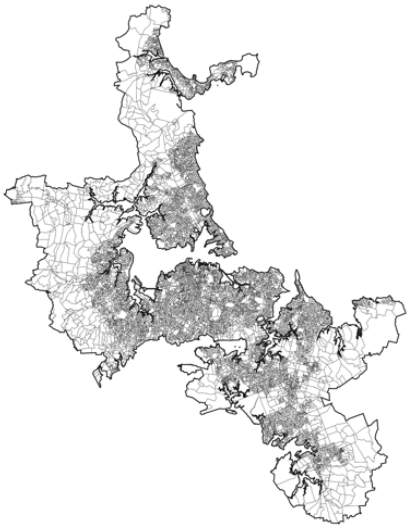

12,017 units.
2018 population: 73 (mean),
108 (median)

(b) Level 2

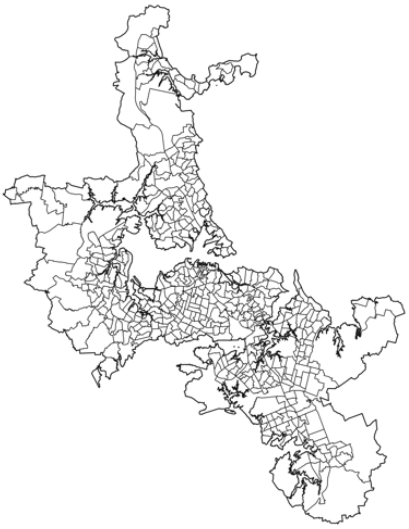

496 units.
2018 population: 2,942 (mean),
3,030 (median)

(c) Level 3

42 units.
2018 population: 34,741 (mean),
32,264 (median)

(d) Level 4

25 units.
2018 population: 58,365 (mean),
53,073 (median)

(e) Level 5

14 units.
2018 population: 104,223 (mean),
97,483 (median)

Note: The maps show the geographic area units used in this paper. The units are organized in a hierarchy of increasing size, where each level's units are constructed by aggregating the previous level's units. Level 1 is set as Statistics New Zealand's meshblock unit, and Level 2 is set as their Statistical Area 2 (SA2) unit. Levels 3-5 are constructed by the author. The number of units in the city and average population per unit (as of the 2018 census) are listed below each map.

B Reduced Form Appendix

B.1 Distance to the Boundary and School

Figure B.1: Boundary Discontinuity Plots

(a) Decile Score

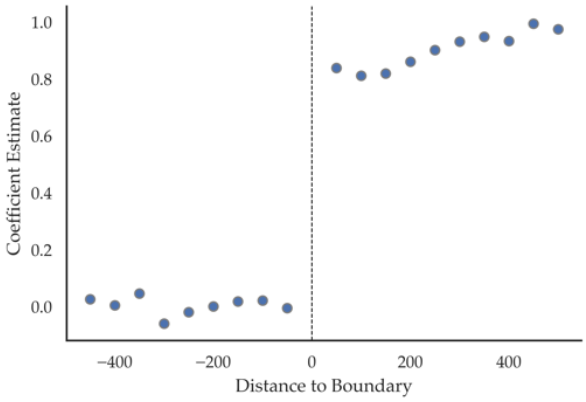

(b) Log Sale Price

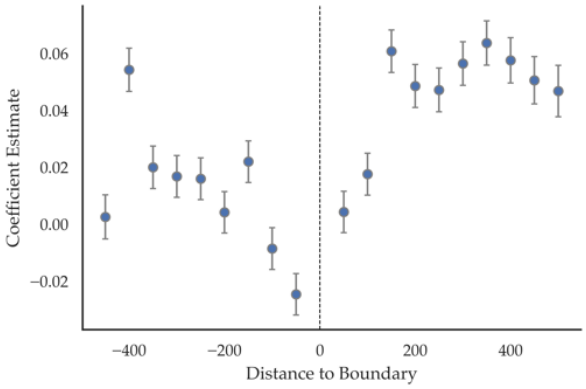

(c) Log Rent Price

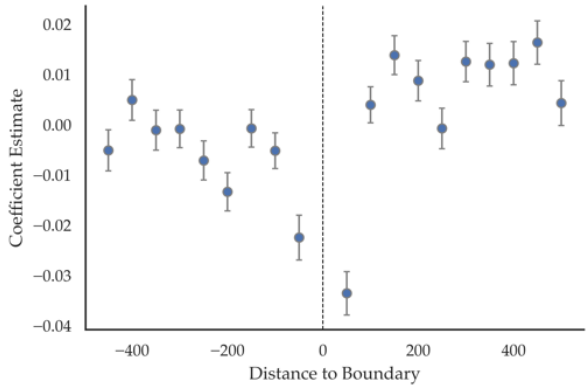

Note: The figures show the mean outcome variable by 50 meter bins of the running variable. The running variable is defined as distance to the boundary, where positive distances indicate the side of the boundary with the higher average decile score. Estimates are calculated using a regression with indicators for each running variable bin as well as year and boundary fixed effects. The distance bin of −500 is the omitted category. 95% confidence intervals are also plotted.

To give a graphical intuition of the BDD identification strategy, I show changes near the boundary in the style of a regression discontinuity plot (Figure B.1).[56] The running variable is the distance to the boundary. I order the boundary sides so that the side with the higher (lower) average decile has positive (negative) values on the running variable. I restrict my analysis to a window of 500 meters on either side of the boundary. In (a), we see that there is a jump in the decile score close to the boundary.[57] Subfigures (b) and (c) show that there also jumps in house prices, both in terms of sales and rents. In

[56] In these plots, I pool boundaries of all school levels into the regression.

[57] The change in decile score at the boundary is positive by construction. Due to the ordering of the boundary sides, the higher quality side is on the right-hand side of the graph.

Figure B.2, I show similar plots for other variables to check for other possible changes at the boundary. One noticeable factor is a spike of sales on areas closest to the boundary (subfigure (c)). This spike occurs on both sides of the boundary (and across many subsamples that I test), so it does not appear to threaten the identification strategy.

Figure B.2: Boundary Discontinuity – Robustness Plots

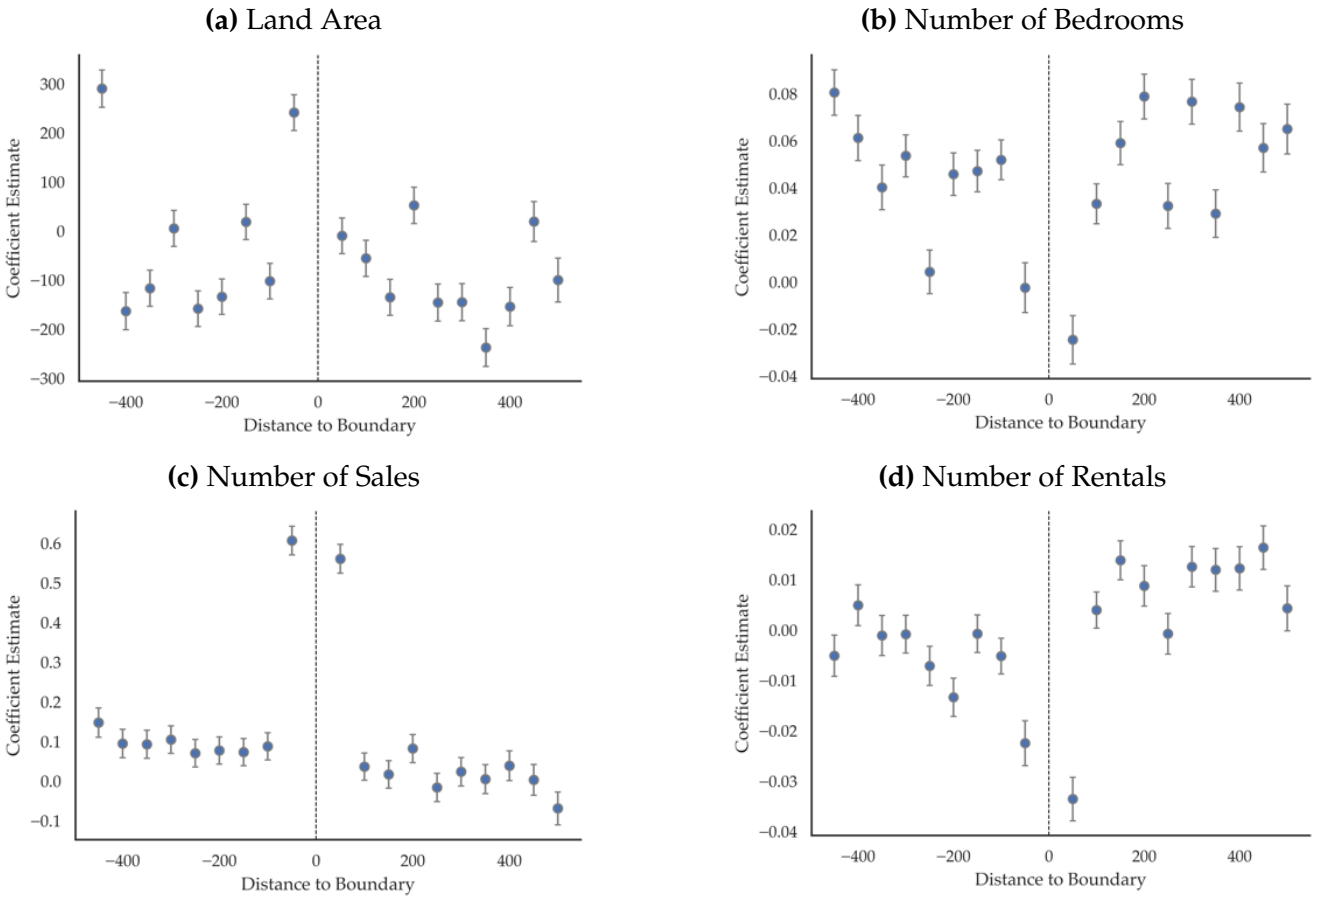

Note: See notes under Figure B.1

Using the family panel, I calculate whether ethnic shares change discontinuously at the boundary. Bayer et al. (2007) find this in their setting (the Bay Area in 1990). Figure B.3 shows that while ethnic shares are changing, it does not appear to be in discontinuous manner. As New Zealand school zones are not binding for admissions, this could point to less extreme sorting at the boundary, as compared to the United States. The school zones will not be explicitly accounted for in the structural model (though they will implicitly affect the school attendance probabilities). Implementing the zones into the structural model would require defining locations in the model to be at a very small geographic unit. This would make the estimation infeasible and/or highly imprecise. Monarrez (2021) finds that 5% of variation in racial and school segregation in the United States is explained by school attendance boundaries. If this effect is lower in New Zealand (as this evidence suggests), then it suggests choosing geographic units larger than school zones can still account for the vast majority of segregation patterns.

Figure B.3: Boundary Discontinuity – Ethnic Shares

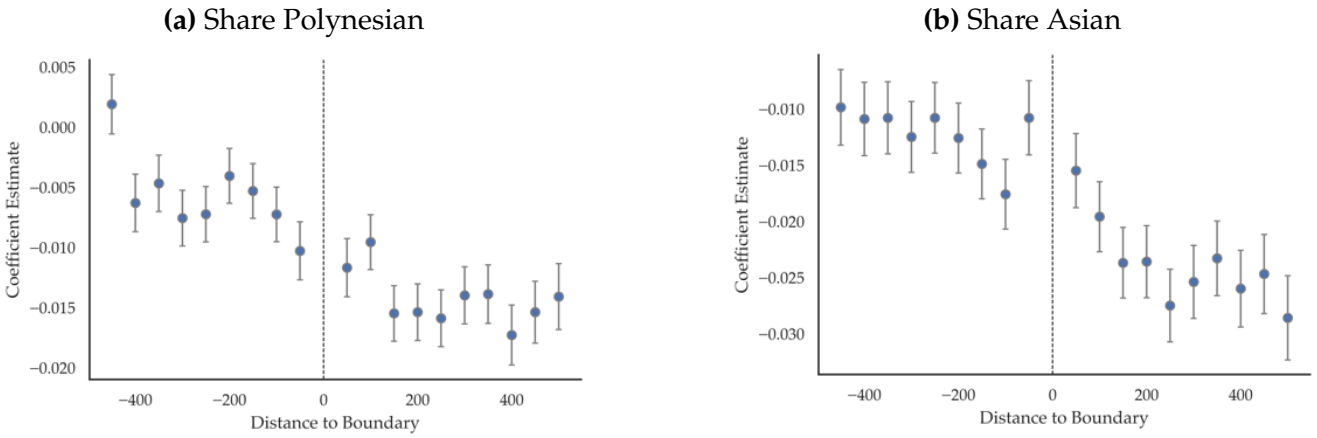

Note: See notes under Figure B.1

One concern is that the BDD necessarily only looks at a narrow geographic area. Moreover, it is an area that is the furthest from the school. Therefore, the people who move into these areas may not fully reflect the general population's preferences. For example, they may place a lower value on schools. Figure B.4 shows how ethnic shares vary with distance from primary schools. There are also similar patterns of changing ethnic shares for intermediate (Figure B.5) and high schools (Figure B.6). Areas very close to low decile schools have a high share of Polynesians, but the white and Asian share increases as distance to the school increases. In contrast, high decile schools have a lower white share further away from the school. By focusing on areas that are far from the school, the BDD analyzes neighborhoods that could differ substantially to the rest of the catchment areas. Therefore, the BDD may have limited external validity, which is a particular concern when studying residential sorting and segregation.

Figure B.4: Ethnic Share by Distance to School (Primary)

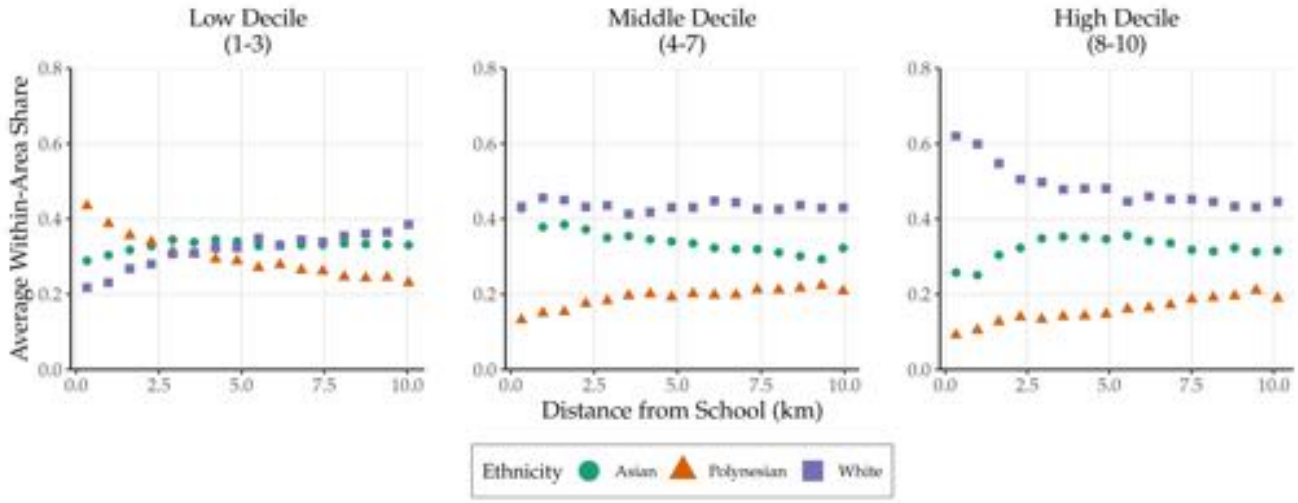

Note: The figure shows a binned scatter plot of the average ethnic share by distance from the school. Ethnic shares are calculated within each Level 2 area (Figure A.16b) using the individual-responses to the 2018 census. Averages are taken over each school-area unit for areas that are within 10 kilometers from its respective school.

Figure B.5: Ethnic Share versus School Distance (Intermediate)

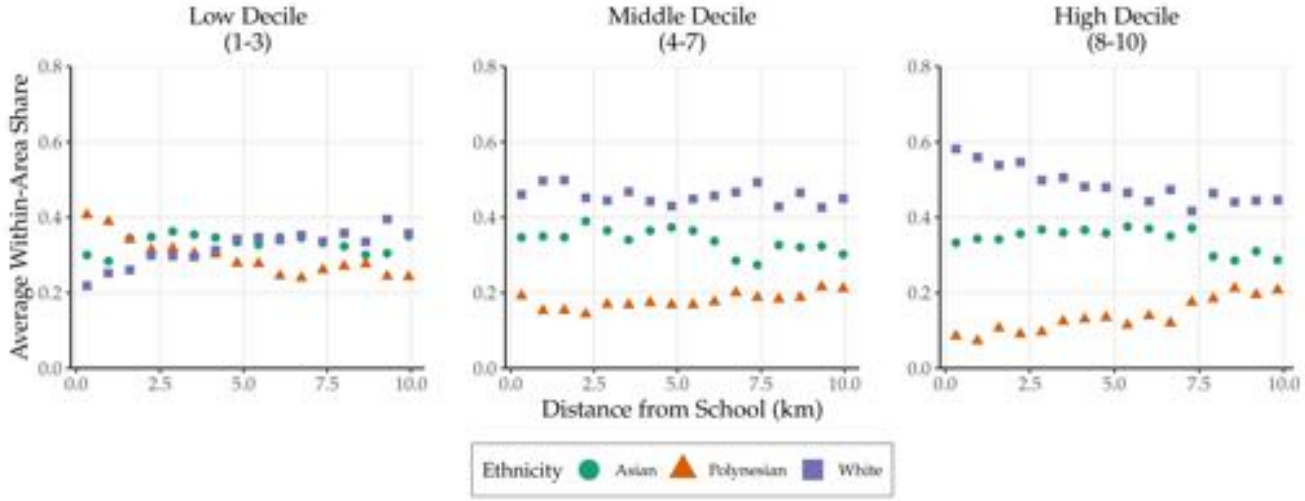

Note: See notes under Figure B.4

Figure B.6: Ethnic Share versus School Distance (High School)

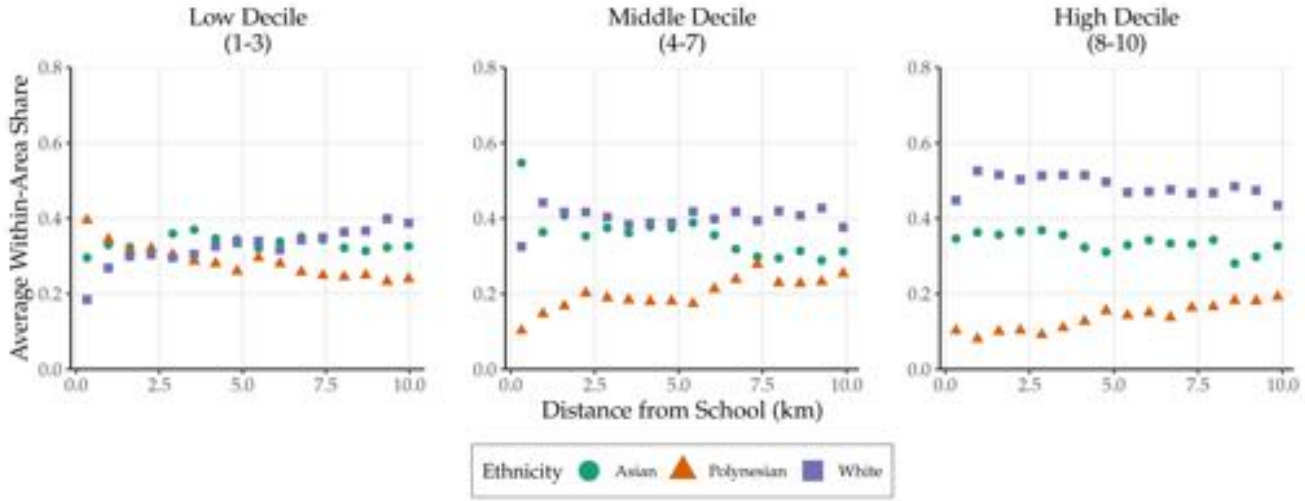

Note: See notes under Figure B.4

B.2 Boundary Construction

In this section, I describe the construction of the boundary sample for the boundary discontinuity analysis.

The New Zealand Ministry of Education provided me with shapefiles for each school zone including historical records to allow for assessing zone changes over time. These files come with the following caveat: "The polygons representing the home zones do not have a high degree of spatial resolution nor vertical alignment integrity. This file should not be used to perform spatial queries".[58]

Figure B.7: School Zone Conversion Examples

(a) Example 1 **(b)** Example 2

Note: The maps show the conversion from the school zones reported in government records to a zone that is composed of meshblocks (the geographic units). In blue, the reported zone is plotted. In orange, the meshblock-based zone is overlaid. The darker regions represent where the two areas overlap.

In using the files, I observe two patterns. First, that the zones largely align with the boundaries of meshblocks. These are geographic units defined by Statistics New Zealand and are analogous to U.S. census blocks (I refer to them as Level 1 units in Figure A.16a). Second, I observe what appears to be slight errors and irregular patterns near the boundary (e.g. borders that are parallel to, but not overlapping with, meshblock boundaries). To account for these errors, and for computational efficiency, I convert each school zone into a zone that is composed of meshblocks. I include a meshblock if more than 50% of its area is covered by the reported zone.[59] Two examples of this process are shown in Figure

[58] Source: https://www.educationcounts.govt.nz/data-services/school-enrolment-zones

[59] This majority-area threshold can generate greater inaccuracies the larger the meshblock. However, note that meshblocks

B.7. While this conversion may cause potential errors in identifying the boundary, it does provide a standardization to the zones and corrects for existing errors in the shapefiles. Moreover, it allows me to express the spatial information in the form of a table, which significantly improves computational efficiency as it minimizes the use of spatial queries.

Figure B.8: Zone Boundary Exclusion Example

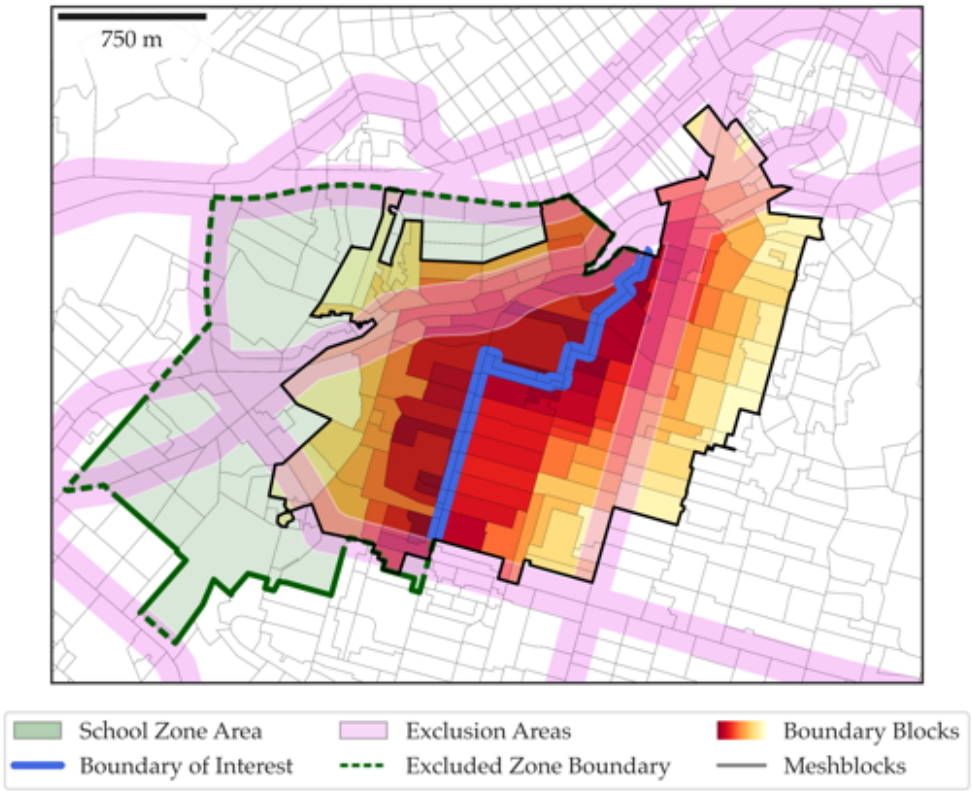

Note: The figure shows how I exclude boundaries and select blocks into the preliminary sample. An example zone is shown in green. The areas to be excluded ("exclusion area") are plotted in pink. Segments of the zone boundary that overlap with the exclusion area are plotted as dashed line. The figure then shows for a segment of the non-excluded boundary (blue segment), the set of blocks that are assigned to it. These are blocks that are less than 500 meters away and not closer to an excluded boundary segment. These boundary blocks are shaded in yellow to red according to their distance to the boundary (darker indicates closer).

The next step is identifying potential boundaries for analysis. As is standard in the literature, I exclude boundaries which may capture changes other than schools. In particular, I exclude any boundaries that are close to highways, train tracks, major roads, as well as the coastline and limits of the urban area ("excluded boundaries").[60] Of the remaining boundaries ("included boundaries"), I find meshblocks whose centroids are less then 500 meters away from the boundary. I remove blocks that are closer to an excluded boundary segment than they are to an included boundary segment. Figure B.8 shows an

are defined by Statistics New Zealand to have approximately the same number of households in each of them. Therefore, meshblocks that are larger are, by definition, more sparsely populated. This means they will also be less relevant for the analysis.

[60] I exclude a buffer around each of these features. I apply a 50 meter buffer to the transportation exclusions and a 100 meter buffer to the coastline/urban area border.

example of this process. Note that I do this for all zones and for all included boundaries.

The final step requires taking into account the uniqueness of the New Zealand system. Typically, catchment areas are mutually exclusive. Therefore a boundary between two catchment areas is easily defined. New Zealand school zones are more irregular than catchment areas in other regions. Here, zones are independently set by schools (rather than being set by a centralized school district). This results in zones that overlap, both within the same level and across school levels (as seen in Figures A.2 and A.3). To ensure that the identification strategy is still valid, I find areas around a school zone boundary such that houses on either side of the boundary are in-zone for the same set of schools, except exactly one school (see Figure B.9). This means that I can attribute the change in house prices to the change of gaining (or losing) guaranteed admissions to exactly one school. Without this, I could then attribute, for example, a change in house prices to a change in primary school quality, when high school quality is also changing across that same boundary. Additionally, these zones can change over time as schools introduce zones or change their existing boundaries. Therefore, I re-determine the set of valid boundaries for each year in the sample to account for these changes.[61]

Figure B.9: Zone Boundary Selection

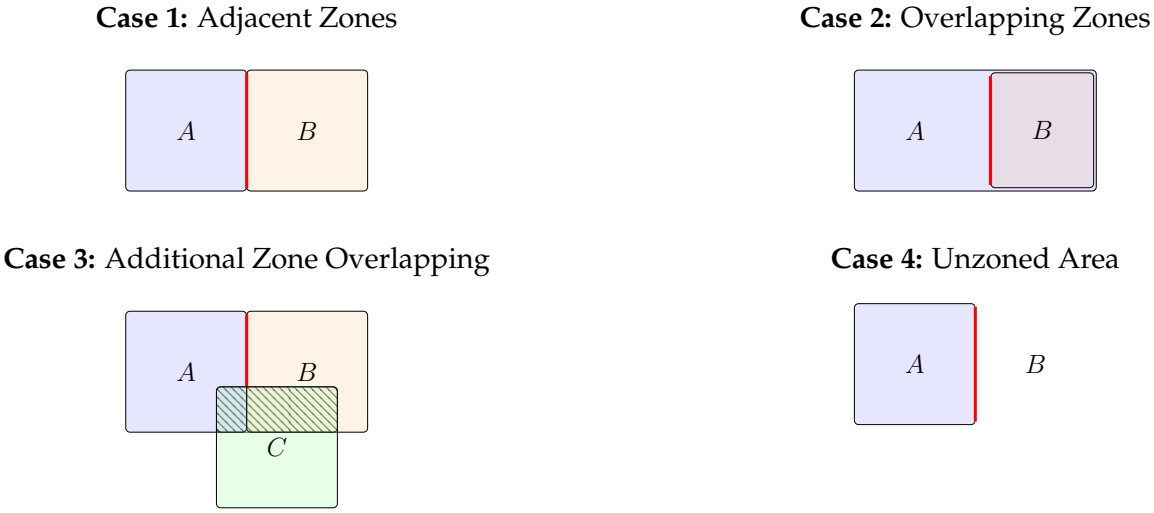

As school zones can overlap, this can make the definition of a "boundary" less clear. I consider three possible situations, which I plot as simplified diagrams in Figure B.9. Consider two zoned schools A and B. If the two zones are adjacent (Case 1), then the boundary is the segment that separates them (highlighted in red). If the two zones overlap (Case 2), then the boundary is still the same. In the diagram, zone B is located within zone A. This means that families in zone B have guaranteed admissions to both schools A and B. However, all that matters for the purposes of identification is whether there is a discontinuous change in expected school quality at the border. The third case is if there is another zone C that overlaps with the boundary (Case 3); then I exclude the border segment that is overlapping with

[61] A boundary is defined only for the years in which there are no changes to the zoning status of houses on either side of the boundary. If the zoning status changes, then I define the boundary as a new boundary with a new fixed effect.

zone C. In addition, when studying the A-B border, I also exclude the areas in A and B that overlap with C (hatched areas). This is to ensure that I only focus on boundaries where one change is occurring. Note that in this diagram it is also possible to study the A-C border and the B-C border separately.

Figure B.10: Comparison of Expected Zone Status

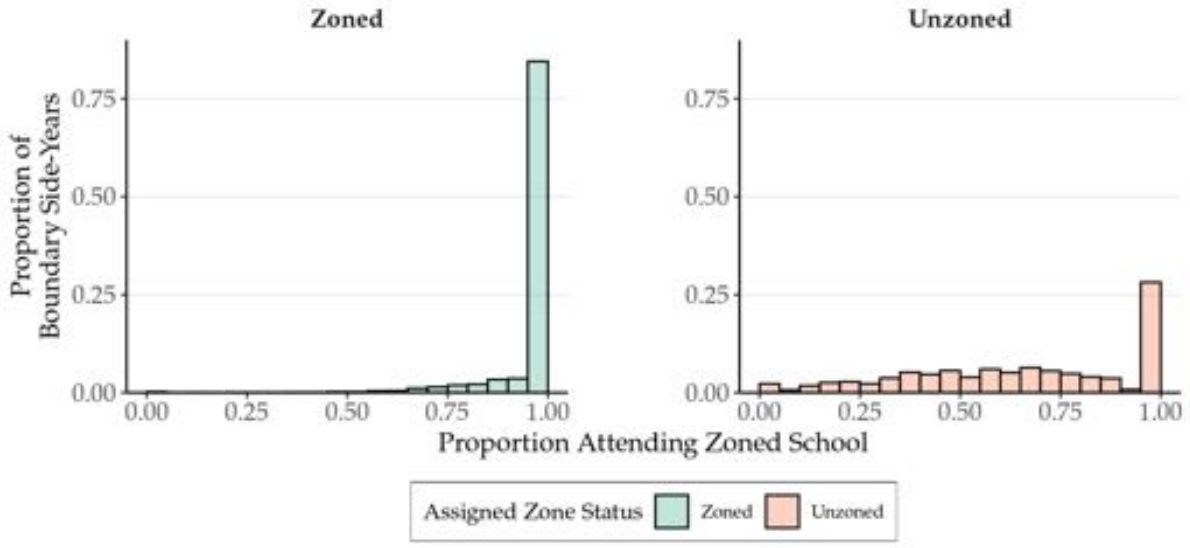

Note: The figure shows the distribution of boundary-side-years based on the proportion living in them who attend a zoned school. The proportion attending a zoned school is shown on the horizontal axis. The proportion over all boundary-side-years (the unit of observation) is shown on the vertical axis. I do this separately by boundary-side-years that are covered by a school zone ("zoned") and by those which are not covered by a school zone ("unzoned").

Another possibility unique to New Zealand is a zoned school A that is adjacent to an unzoned school B (Case 4 in Figure B.9). People living in these unzoned areas can still apply to any school; however, they do not receive guaranteed admissions to the zoned schools. For my analysis, I exclude boundaries where one side of the boundary is unzoned due to the potential for selection bias. Using information about children's home and school from the family panel, I determine what proportion of students living in a boundary-side are attending a zoned school. I calculate this statistic for every boundary-side-year in my sample and plot the distribution in Figure B.10. Unsurprisingly, I find that boundary-sides that are covered by a school zone tend to have almost all their students attending a zoned school (left panel). However, I find that unzoned boundary-sides still have a high proportion of their students attending zoned schools (right panel). This suggests that being unzoned, at least close to the boundary, is typically not a binding treatment. School attendance patterns in these areas appear to be driven more by parent preferences than exogenous school assignment policies. This induces a correlation between (selected) school quality and the error term, which may change discontinuously at the border and would therefore not be accounted for by the boundary fixed effects. Additionally, the effect of being in a zoned area could have an additional treatment effect (e.g. if zoning is viewed in of itself as more prestigious), which may also bias the estimates of school quality valuation. To avoid these issues, I only focus on the analysis where both sides of the boundary are zoned. This is also makes it more comparable to studies of other settings, which often have mutually exclusive catchment areas that span the entire city.

B.3 BDD School Quality Assignment

It is important to note that school zones change the probability of admissions (conditional on applying), but families in New Zealand always have a choice in which school their child will attend. This raises the concern that school zones may not be generating a meaningful difference in school attendance. In other settings, catchment areas are typically binding and therefore the assignment of houses to schools is relatively straightforward. The New Zealand setting can be considered closer to a fuzzy regression discontinuity design. The identification strategy still follows as the exogenous variation will be of a discontinuous change in the *expected* school quality.

The expected school quality can be calculated as the average decile score, weighted by the probability of attending each school (in a given boundary-side-year). This can be interpreted as the ex ante school quality a family would expect to enjoy if they moved into the area. However, this method requires attendance probabilities to be estimated for all boundary side-years (as zones can change over time, these probabilities need to be re-calculated for each year). This is empirically challenging as the BDD requires defining a narrow area around a boundary, which means that these probabilities are estimated from potentially small samples.

Figure B.11: Comparison of Expected School Quality

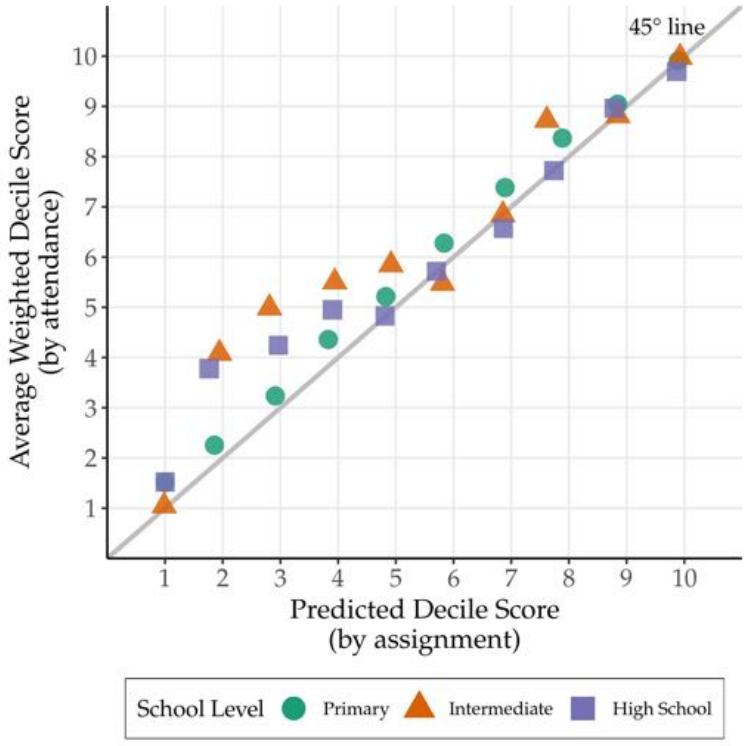

Note: The figure shows a comparison between the expected decile score using a simple assignment rule versus expected decile score using attendance probability weights from the empirical distribution. Each unit is a boundary-side-year, which have been binned by their predicted (assignment-based) decile score to calculate the mean observed (attendance-based) decile score.

To circumvent this, I use an assignment rule instead of weighting by probabilities. I assign each home to a school using a simple rule: school quality is calculated as the average quality for all schools that the house is in-zone for; if it is not in-zone for any school, then assign the quality of the nearest unzoned school.

For areas with large enough samples, I am able to calculate the actual attendance probabilities. There are areas that have large enough populations on both sides of the boundary. I use this to check whether the expected school quality using attendance weights matches the expected quality using my assignment rule. Figure B.11 shows that these align for the most part. One exception is intermediate and high schools between deciles 2 to 4. It appears in these areas that families are sending their children to much higher decile scores that one would expect. I find that this is largely driven by the unzoned areas, which will be dropped for the reasons explained above.

Figure B.12: Changes in Attendance Probability

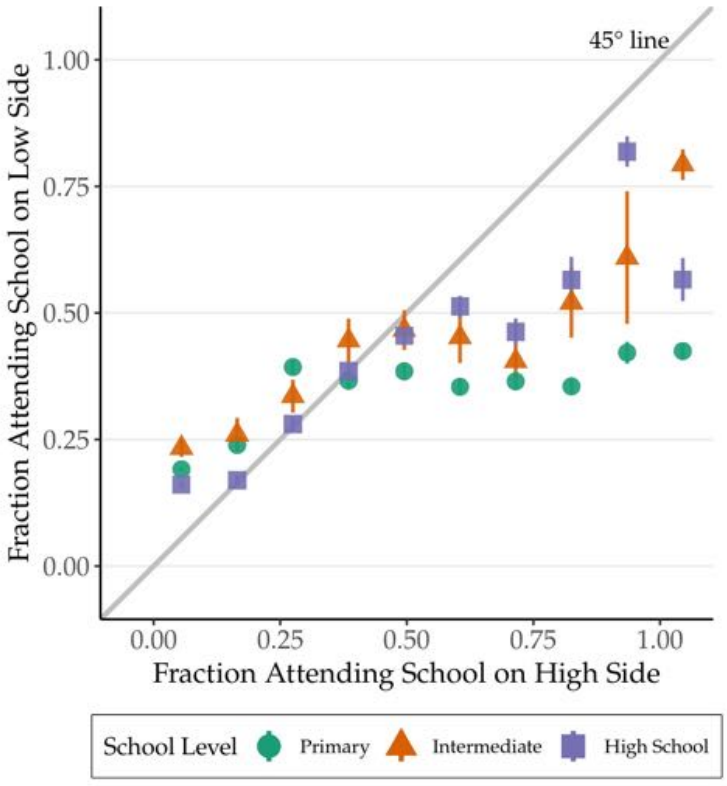

Note: The figure compares how attendance probabilities for a school vary on each side of a boundary-year. Sides of a boundary are categorized as "low" and "high" based on whichever has the higher average decile score. The figure plots a binned scatterplot by binning attendance probabilities on the high boundary side. The mean attendance probability (for the same school) on the low side is plotted on the vertical axis. The 45° line indicates equal attendance probabilities on both sides. Points above the 45° line indicate schools that are more likely to be attended on the low side than the high side. Points below the 45° line indicate schools that are more likely to be attended on the high side than the low side.

Using these probabilities, I also check to see if there is a discontinuous change in the attendance probability at the the zone boundary. As in the boundary discontinuity plots of Figure B.1, I categorize the

boundaries side into "high" and "low" according to which side has the higher average decile score. For each side of the boundary (in each year), I calculate the probability of attending each school. This gives data at the boundary-side-year-school level. I then compare the probability of attending a school m at the low side of a boundary-year in comparison to the probability of attending the same school m but on the high side of the same boundary-year. This is shown as a binned scatterplot in Figure B.12. I overlay the 45° line, which represents equal attendance probabilities on both sides of the boundary-year. Points that deviate from the 45° line are interpreted as a discontinuity in the probability of attendance. We see that while the attendance probability on the two sides are correlated, there appears to be a discontinuity, especially at the extreme points. In particular, for schools that have a more-likely-than-not probability (>50%) of attendance on the high side, there is a less than 50% chance of attending these schools on the low side.

B.4 BDD Robustness Checks and Benchmarking

In Tables B.1 and B.2, I conduct further robustness checks. I check for different restrictions on the distance from the boundary and find that the results hold even when I limit to houses 250 meters away from the boundary. Dropping locations very close to the boundary results in larger WTP estimates, which could be driven by sorting patterns. As schools can change their zones, it is possible that those who value school quality the least are willing to live near a boundary that could potentially change in the future.[62] Another concern is that boundaries are defined over narrow areas, which could result in focusing on sparsely populated areas. I restrict to areas with at least 50 families living on both sides of the boundary (calculated via the family panel) and find the same results.

To put the results into context, other papers in the literature find that a 1 standard deviation increase in elementary school quality causes a 3-4% increase in house prices (Black and Machin, 2011). For the New Zealand setting, "a standard deviation increase in school quality" is best thought of as a one standard deviation increase in the underlying (continuous) measure used to calculate the decile score (Figure A.4). I calculate that such an increase corresponds to an increase of approximately 3 decile scores. Therefore, a comparable result is that I find that a one standard deviation increase in elementary school quality causes a 2.6% increase in house prices. To be clear, this is not a perfect comparison. The decile score is a relative measure (a ranking), where the interpretation of standard deviation increase is not as clear as in the case of test scores. However, this does suggest that the valuation of quality (as measured by a BDD) is similar as compared to other contexts. My results are qualitatively similar to Caetano (2019), who uses a structural model to estimate that parents most value high school, followed by elementary schools, and middle school least.[63]

[62] Another reason could be due to errors in the geographic boundary information, as discussed in Appendix B.2, that results in measurement error close to the boundary.

[63] Caetano (2019) finds a WTP (per child per year) of 37.1%, 21.2%, and 41.2% for a one standard deviation increase in elementary, middle, and high school quality, respectively. These magnitudes are substantially larger than others in the literature.

Table B.1: Robustness Checks (Sales)

	(1) **Between 0−250m**	(2) **Between 50−500m**	(3) **Between 50−250m**	(4) **Boundary Year FE**	(5) **Attendance-Based Decile**
School Quality					
High School	0.0336**	0.0458***	0.0477**	0.0548**	0.0709***
	(0.0147)	(0.0133)	(0.0178)	(0.0220)	(0.0258)
Intermediate	0.0006	−0.0103	−0.0063	−0.002	−0.0559*
	(0.0115)	(0.0107)	(0.0087)	(0.0123)	(0.0315)
Primary	0.0091**	0.0069**	0.0084**	0.0082**	0.0060
	(0.0039)	(0.0031)	(0.0037)	(0.0034)	(0.0058)
House Chars.	Yes	Yes	Yes	Yes	Yes
Ethnic Shares	Yes	Yes	Yes	Yes	Yes
N	75,795	111,084	58,233	128,646	105,066
R^2	0.822	0.817	0.824	0.825	0.82
Clusters	39	39	39	39	39

	(6) **Years 2009−2014**	(7) **Years 2015−2019**	(8) **Populated Areas**	(9) **Level 4 Cluster**	(10) **Boundary Cluster**
School Quality					
High School	0.0608**	0.0442**	0.0383**	0.0350***	0.0350***
	(0.0293)	(0.0182)	(0.0174)	(0.0121)	(0.0132)
Intermediate	−0.0015	−0.0031	−0.0096	−0.0024	−0.0024
	(0.0096)	(0.0209)	(0.011)	(0.0148)	(0.0110)
Primary	0.0107***	0.0045	0.0072**	0.0067**	0.0067*
	(0.0037)	(0.0042)	(0.0035)	(0.0028)	(0.0036)
House Chars.	Yes	Yes	Yes	Yes	Yes
Ethnic Shares	Yes	Yes	Yes	Yes	Yes
N	69,918	58,728	103,365	128,646	128,646
R^2	0.818	0.794	0.818	0.816	0.816
Clusters	37	39	39	23	452

Note: The table shows extensions to the regressions shown in Table 1a (see its notes for further details). Columns (1)-(3) subset to houses within the specified distance from the boundary (baseline uses a 500m window). Column (4) uses a combined boundary-year fixed effects. Column (5) uses the decile calculated based on empirical attendance probabilities (if available). Columns (6)-(7) subset the years. Column (8) subsets to areas with more than 50 families living in each boundary side. Column (9) uses Level 4 area (Figure A.16d) as the clustering variable (baseline uses Level 3). Column (10) clusters at the boundary level. Significance stars thresholds: $^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$. In accordance with Statistics New Zealand confidentiality rules, the number of observations has been randomly rounded to a multiple of 3 (RR3).

Table B.2: Robustness Checks (Rents)

	(1) **Between 0−250m**	(2) **Between 50−500m**	(3) **Between 50−250m**	(4) **Boundary Year FE**	(5) **Attendance-Based Decile**
School Quality					
High School	0.0224***	0.0377***	0.0338***	0.0353**	0.0238
	(0.0068)	(0.0130)	(0.0111)	(0.0151)	(0.0200)
Intermediate	0.0060	0.0068	0.0040	0.0068	−0.0115
	(0.0050)	(0.0074)	(0.0059)	(0.0076)	(0.0310)
Primary	0.0008	0.0028	0.0058*	0.0009	0.0010
	(0.0032)	(0.0038)	(0.0032)	(0.0042)	(0.0057)
House Chars.	Yes	Yes	Yes	Yes	Yes
Ethnic Shares	Yes	Yes	Yes	Yes	Yes
N	178,650	292,284	157,527	313,404	237,456
R^2	0.55	0.552	0.551	0.567	0.546
Clusters	39	39	39	39	39

	(6) **Years 2009−2014**	(7) **Years 2015−2019**	(8) **Populated Areas**	(9) **Level 4 Cluster**	(10) **Boundary Cluster**
School Quality					
High School	0.0514**	0.0239**	0.0388**	0.0281***	0.0281***
	(0.0210)	(0.0114)	(0.016)	(0.0067)	(0.0103)
Intermediate	0.0049	0.0127	0.0062	0.0086	0.0086
	(0.0100)	(0.0090)	(0.007)	(0.0082)	(0.0067)
Primary	0.0020	−0.0008	0.0003	0.0002	0.0002
	(0.0053)	(0.0035)	(0.0004)	(0.0045)	(0.0022)
House Chars.	Yes	Yes	Yes	Yes	Yes
Ethnic Shares	Yes	Yes	Yes	Yes	Yes
N	157,803	155,601	233,574	313,404	313,404
R^2	0.568	0.537	0.542	0.551	0.551
Clusters	38	39	39	23	508

Note: The table shows extensions to the regressions shown in Table 1a (see its notes for further details). Columns (1)-(3) subset to houses within the specified distance from the boundary (baseline uses a 500m window). Column (4) uses a combined boundary-year fixed effects. Column (5) uses the decile calculated based on empirical attendance probabilities (if available). Columns (6)-(7) subset the years. Column (8) subsets to areas with more than 50 families living in each boundary side. Column (9) uses Level 4 area (Figure A.16d) as the clustering variable (baseline uses Level 3). Column (10) clusters at the boundary level. Significance stars thresholds: $^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$. In accordance with Statistics New Zealand confidentiality rules, the number of observations has been randomly rounded to a multiple of 3 (RR3).

B.5 Additional Moving Profile Analysis

Figure B.13 shows the family's moving probability relative to the mother's age. In the moving profile analysis, I always control for mother's age and normalize the results to be relative to the moving probability for a family with a 35 year-old mother.

Figure B.13: Moving Probabilities – Mother's Age

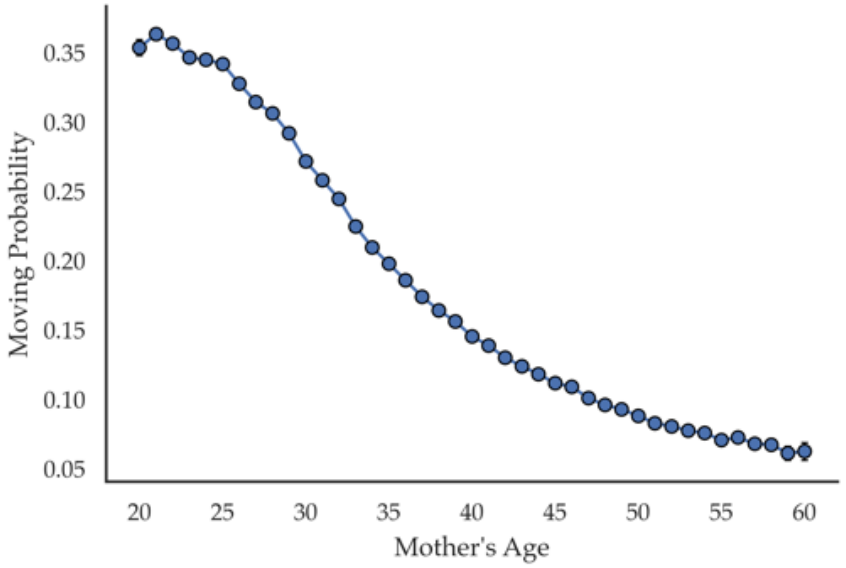

Note: The figure shows the moving probability for the family given the mother's age. Estimates are from a regression of a moving indicator on dummies for each age of the mother. 95% confidence intervals are plotted.

In Figure B.14, I show that patterns seen in Figure 4 generalize to families with multiple children. I find that the intertemporal moving profile differs between the eldest and youngest child. Notably, families are far less likely to move at the time of youngest child's birth as compared to the time of the eldest child's birth. This suggests that moves when children are born are not simply explained by, for example, moves to larger houses to accommodate the larger family. It is however, once again, consistent with dynamic behavior. If parents are forward-looking and account for their future needs when their eldest child is born, then this lowers the need to re-adjust their location when their other children are born. This figure also emphasizes that a structural model needs to take into different family sizes and the children's birth order when considering the family's moving choice.

Finally, to show that schools (and not other factors) are driving the effects seen in Figure 4, I leverage cut-offs in school entry rules to identify plausibly exogenous variation in the child's grade. Traditionally in New Zealand, children enroll in primary school close to their 5^{th} birthday. Children who enroll in the first half of the year are placed in first grade and begin their education at age 5. Children who enroll in the later half of the year are placed in a "grade zero" cohort that accommodates them until they can enter first grade at the start of the following year. These children therefore begin first grade at age 6. As a result, two children born in the same calendar year may be in different school grades due to their

Figure B.14: Moving Probabilities – Child's Grade and Birth Order

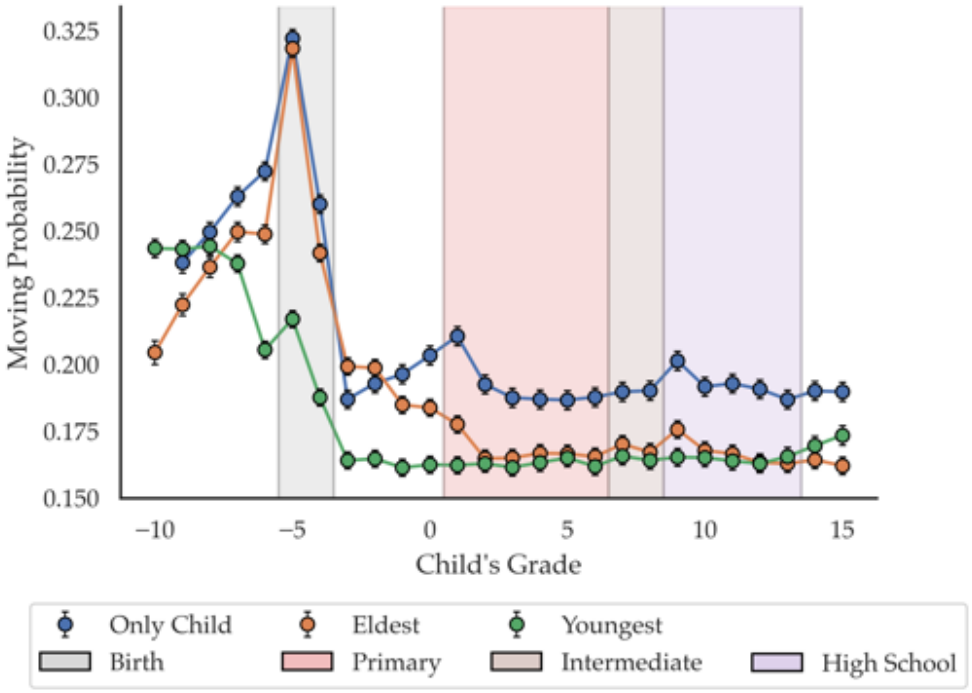

Note: The figure shows the moving probability given the grade and birth order of the child. I regress a moving indicator on dummies for each grade of the child (interacted with birth order) and dummies for each age of the mother. Each point is equal to the grade-order-specific coefficient plus the coefficient for age 35 mother. 95% confidence intervals are plotted. Middle children are excluded from the regression. Depending on the age they start school, the child's birth is represented at either grade −5 or −4. Shaded regions indicate the ranges for each school level: primary (grades 1-6), intermediate (grades 7-8), and high school (grades 9-13).

month of birth. As Figure B.15 shows, those born in July or later tend to start first grade at age 6, while those born before March tend to start at age 5.[64]

I estimate the moving probabilities by age (rather than grade) of the child for families of two cohorts: children born between January and March (age 5 starters), and children born between July to December (age 6 starters). Figure B.16 shows that the two cohorts have similar moving profiles before children start school, which is expected as they are born in the same year.[65] However, after they start school, we see the age 5 starters have a moving profile that is shifted by one period forward as compared to the age 6 starters. In particular, we see the same spikes as in Figure 4, however, their timing reflects that the determinant is the *grade* of the child rather than their *age*.

[64]Schools can set their own entry cut-off dates, which tend to range between April and July. Parents can also choose when their child begins school (until age 6, after which school is compulsory). For my analysis, I exclude those born between April to June due to potential selection bias.

[65]For the July-December cohort, one noticeable difference is that their families are more likely to move at age 1. This is likely explained by the fact that they are born in the later half of the year, while I determine residence at an annual basis based on the most lived house within each year. Therefore, the effect that a new child has on moving propensity likely would not be fully captured for this cohort until the year following their birth.

Figure B.15: Age at First Grade by Month of Birth

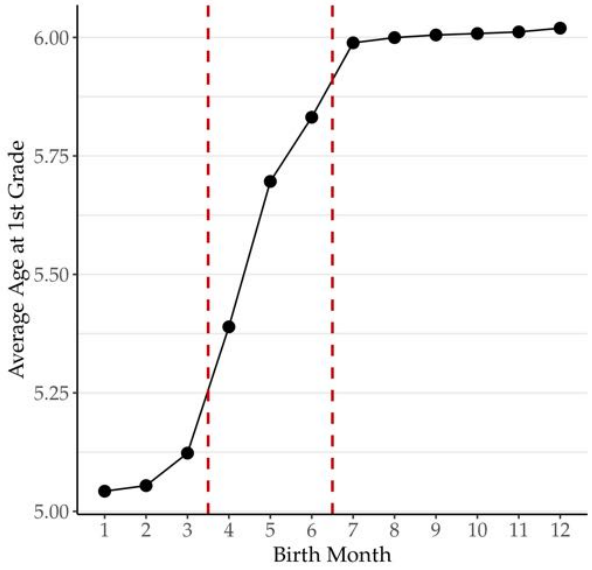

Note: The figure shows the mean age at first grade by the month of birth for all children in the sample data. Age is calculated as year enrolled in first grade minus year of birth. Vertical lines represent the range of dates (1 April to 30 June) where schools typically set their entry cut-offs.

Figure B.16: Moving Probabilities – Child's Age

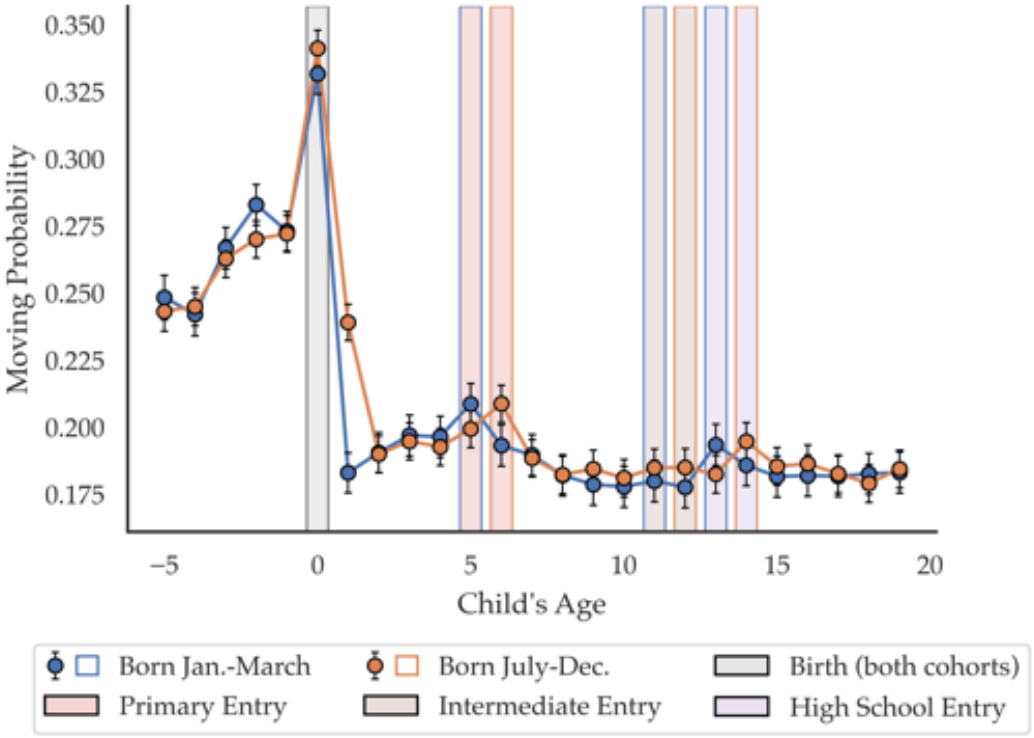

Note: The figure shows the moving probability given the age of the child. I regress a moving indicator on dummies for each age of the child (interacted with a cohort indicator) and dummies for each age of the mother. Each point is equal to the age-cohort-specific coefficient plus the coefficient for age 35 mother. 95% confidence intervals are plotted. The cohort indicator represents being in born between January–March or July–December (those born April–June are excluded). The first cohort begins school at age 5, while the second cohort begins school at age 6. Shaded regions indicate child's birth (age 0) and the start of each school level for the respective cohorts.

C Model Appendix

C.1 Model Details

C.1.1 Children, Entry, and Exit

The model focuses on the choices of families with children. This means I need to specify: (i) the time periods for which this model represents a family's decision-making, and (ii) how the family composition changes as children are born.

On the first point, my goal is to study the family's choices for the periods where children are most likely to play an important role. This means I restrict the decision window for each family to be a few years before a family has their first child up to a few years after their youngest child finishes high school. This simplification is to keep the focus on the residential choice of families rather than groups such as childless college students or retirees. These groups may have very different preferences from families, which would not be well captured by this model.[66]

On the second point, I simplify the problem by abstracting away from how families decide when to have children (and how many to children to have). I assume that the total number of children each family will have is exogenous and fully anticipated by the parents. However, I introduce randomness by having the timing of each subsequent birth to be drawn from an exogenous distribution. While the timing of subsequent births is random, I model the parents as fully anticipating the birth of a child by some number of periods. Figure C.1 illustrates how this process plays out for a family with two children. In reality, the true decision-making around having children is far more complex; however these simplifications still capture the main patterns of family composition observed in the data.

Next, I specify the above assumptions using the model notation. Each family i has $C_i \in \mathbb{N}^+$ children. A child $c \in C_i$ is born in period T^b_{ic} and therefore is of age $a_{itc} = t - T^b_{ic}$ at time t. C_i is fully anticipated and exogenous. The probability that child c will be born in period t is based on a distribution B_ψ:

$$Pr(T^b_{ic} = t) = B_\psi(C_i, C_{it}, a_{it,c-1}) \tag{C.1}$$

where C_{it} is the number of children in family i at time t, and $a_{it,c-1}$ is the age of the previous child at time t. Note that this distribution varies by type ψ to allow for different family compositions across types. For a child born in period T^b_{ic}, the family will know $\underline{b}$ periods prior that the child will be born, i.e. they know in period $T^b_{ic} - \underline{b}$ that their child c will be born in period T^b_{ic}.

In a given period, there will be N_t agents in the economy. Families enter the economy under two possible

[66] As Figure C.6 shows, the spatial distribution of parents can look very different to these (childless) younger and older groups, suggesting fundamentally different location preferences.

scenarios. They either enter when their eldest child's ($c = 1$) birth is known, which is at period $T_i^0 = T_{i1}^b - \underline{b}$. Alternatively, they can enter exogenously through migration (where their children can be of any age upon entry). When families enter, their initial location will be assigned exogenously. Agents also exit the model once their youngest child ($c = C_i$) reaches adulthood, which will be defined as $\bar{b}$ periods after birth. This is denoted as $T_i^1 = T_{i,C_i}^b + \bar{b}$.

Figure C.1: Child Birth Process

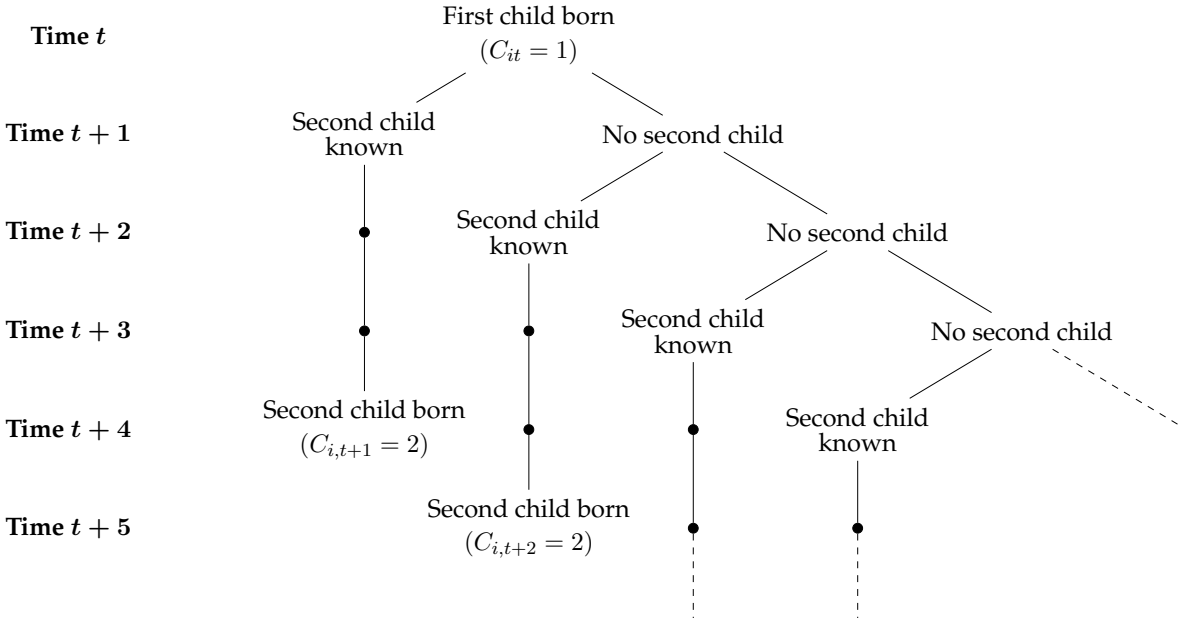

Note: The diagram illustrates the family composition transition for a two-child family ($C_i = 2$) who has their first child born in period t and do not yet know the timing of their second child's birth. The diagram shows the possible options at each stage, with an announcement lag of three periods ($\underline{b} = 3$). After the second child is born, they know with certainty that there will be no more children.

C.1.2 Intraperiod Timing

In period t, the family chooses their action d_t as a function of the state $\mathbf{s}_{it} = (k_{it}, w_t, \eta_t, \varepsilon_{it})$. However, k_{it} is actually a vector of $t-1$ variables, such as home location ($h_{i,t-1}$) and tenure ($\tau_{i,t-1}$). These $t-1$ variables represent the family's *current* living situation at the start of period t and are simply expressed as k_{it} for notational convenience.

The timing in the model works as follows:

1. Period t starts. $k_{it} = (\psi_i, h_{i,t-1}, \tau_{i,t-1}, m_{i,t-1}, a_{i,t-1}, g_{i,t-1}, C_{i,t-1})$ and $\omega_t = (w_t, \eta_t)$ are known at the start of the period.
2. ε_{it} shocks are realized and observed by the agents but not the econometrician.
3. Agents choose their action (d_t) based on $\mathbf{s}_{it}$ and their expectations of how these variables will evolve over time.
4. Flow utilities for period t are realized.
5. Period t ends.
6. State transition occurs:
   - k_{it} transitions to $k_{i,t+1}$, e.g. families move into their new locations (h_{it}), children move into the next grade (g_{it}), and the children's school assignment is realized (m_{it}).
   - Market variables transition from ω_t to ω_{t+1}.
7. Period $t+1$ starts. Now, $k_{i,t+1}$ and ω_{t+1} are known.

C.1.3 Observable Flow Utility Components

Table C.1 provides a summary of the components of $\bar{u}(j, k_{it}, w_t)$. In this section, I provide the specific parameterizations of each component to complement the descriptions given in section 4.3.

The school quality (SQ) experienced by the family is:

$$\begin{aligned} SQ(k_{i,t+1}, w_t) = \Big(& \theta_\psi^{Prim} \cdot \mathbb{1}\{g_{it} \in [1,6]\} + \\ & \theta_\psi^{Int} \cdot \mathbb{1}\{g_{it} \in [7,8]\} + \\ & \theta_\psi^{HS} \cdot \mathbb{1}\{g_{it} \in [9,13]\} \Big) \times q_t(m_{it}) \end{aligned} \tag{C.2}$$

where $q_t(m_{it})$, a component of w_t, is the school quality measure of the school m_{it} at time t. The parameters $\theta_\psi^{Prim}, \theta_\psi^{Int}, \theta_\psi^{HS}$ capture the valuation of quality at the primary, intermediate, and high school level,

respectively (all of which can differ by type ψ). The relevant parameter is used depending on the child's current school level.[67] Note that per the notation in section 4.2, g_{it} and m_{it} are components of $k_{i,t+1}$.

The school moving cost (SMC) is expressed as:

$$\begin{aligned} SMC\left(k_{i,t+1}, m_{i,t-1}\right) = \Big(& \upsilon_{\psi}^{Prim} \cdot \mathbb{1}\left\{g_{it} \in [2,6]\right\} + \\ & \upsilon_{\psi}^{Int} \cdot \mathbb{1}\left\{g_{it} \in \{8\}\right\} + \\ & \upsilon_{\psi}^{HS} \cdot \mathbb{1}\left\{g_{it} \in [10,13]\right\} \Big) \times \mathbb{1}\left\{m_{it} \neq m_{i,t-1}\right\} \end{aligned} \tag{C.3}$$

As the expression shows, this cost is paid as a child changes schools ($m_{it} \neq m_{i,t-1}$), but not if the child begins a new school level (where $g_{it} = 1, 7, 9$ are the starting grades for primary, intermediate, and high school, respectively). Consistent with the other school variables, there is a separate school moving cost parameter associated with each school level and family type. We should expect all these parameters ($\upsilon_{\psi}^{Prim}, \upsilon_{\psi}^{Int}, \upsilon_{\psi}^{HS}$) to be negative.

The distance to school component, represented by $DS(k_{i,t+1})$, is calculated using the distance between the home location (h_{it}) and the assigned school (m_{it}). As in Equation (C.2), I include separate coefficients on distance for each school level and family type.

$$\begin{aligned} DS(k_{i,t+1}) = \Big(& \kappa_{\psi}^{Prim} \cdot \mathbb{1}\left\{g_{it} \in [1,6]\right\} + \\ & \kappa_{\psi}^{Int} \cdot \mathbb{1}\left\{g_{it} \in [7,8]\right\} + \\ & \kappa_{\psi}^{HS} \cdot \mathbb{1}\left\{g_{it} \in [9,13]\right\} \Big) \times \text{dist}(h_{it}, m_{it}) \end{aligned} \tag{C.4}$$

For the outside option school, I set the utility from school quality (SQ) and distance (DS) to be zero. There is still a moving cost associated with changing schools between the outside option and the public schools. Note that the outside option school is different to the outside option location. Children living inside the city can attend the outside option *school* (e.g. a private school). In contrast, the outside option *location* has no explicit school payoff associated with it (see Assumption 5). The outside option component to the school utility, $OOS(k_{i,t+1})$, has a different (fixed) value by school level and family type in the same way as Equation (C.2). Let $m_{it} = 0$ represent attending the outside option school.

$$\begin{aligned} OOS(k_{i,t+1}) = \Big(& o_{\psi}^{Prim} \cdot \mathbb{1}\left\{g_{it} \in [1,6]\right\} + \\ & o_{\psi}^{Int} \cdot \mathbb{1}\left\{g_{it} \in [7,8]\right\} + \\ & o_{\psi}^{HS} \cdot \mathbb{1}\left\{g_{it} \in [9,13]\right\} \Big) \times \mathbb{1}\left\{m_{it} = 0\right\} \end{aligned} \tag{C.5}$$

[67] Recall that primary school is grades 1-6, intermediate is grades 7-8, and high school is grades 8-13.

Finally, there is a moving cost, which is parameterized as:

$$MC\left(j, k_{it}\right) = \begin{cases} \upsilon_{\psi}^{F} + \upsilon_{\psi}^{D} \cdot \text{dist}\left(j, h_{i,t-1}\right) & \text{if } j \in \mathcal{L} \text{ and } h_{i,t-1} \in \mathcal{L} \\ \upsilon_{\psi}^{O} & \text{if } j = 0 \text{ or } h_{i,t-1} = 0 \\ 0 & \text{if } j = stay \end{cases} \tag{C.6}$$

where υ_{ψ}^{F} is a fixed cost to moving and υ_{ψ}^{D} is a per kilometer moving cost. These are paid for moves within the city. For moves between the city and the outside option, there is a fixed cost of υ_{ψ}^{O}. Families who stay in their current home do not pay a moving cost. If there is a disutility to moving (as we would expect), then the parameters $\upsilon_{\psi}^{F}, \upsilon_{\psi}^{D}, \upsilon_{\psi}^{O}$ will be negative.

Table C.1: Observable Flow Utility Components

Category	**Component**	**Symbol**	**Interpretation**
School	School Quality	SQ	Measure of school quality experienced by the family in location l at time t
	Distance	DS	Distance to school from location l
	School Moving Cost	SMC	Cost of a child changing their school within the same school level
	Outside option	OOS	Value of outside option (non-public) school
Neighborhood	House Prices	r	Price index for housing for location l at time t
	Type Shares	TS^{φ}	The share of the population in location l at time t that are of type $\varphi \in \Psi$
Moving and Location	Location Effects	δ	Time-invariant characteristics for location l
	Tenure	τ	Number of years in current home
	Moving Cost	MC	Cost of moving from current location to location l

C.1.4 Child-Related Transition Processes

Recall that the number of children C_{it} transitions as described in section C.1.1 and Equation (C.1).

For the school variables, I distinguish between the age and the grade of the child, as these do not map one-to-one. In New Zealand, children born between January and April tend to be age 5 in the first grade, while those born May to December tend to be age 6 (Figure B.15). Therefore, assuming the month of birth is exogenous, I express the grade transition as:

$$g_{itc} = \begin{cases} a^* - 4 & \text{with probability } 1/3 \text{ if } a_{itc} = a^* \\ a^* - 5 & \text{with probability } 2/3 \text{ if } a_{itc} = a^* \\ g_{i,t-1,c} + 1 & \text{if } a_{itc} > a^* \end{cases} \qquad \forall c \in \{1, ..., C_{it}\}$$

At age a^*, children are assigned randomly to an initial grade cohort.[68] After that, their grade also increases by one each year. There could be a concern that some children may not progress by one grade each year (e.g. if they are held back or skip a grade). However, I will simply assume this away as it is empirically infrequent.[69]

The child's school assignment, m_{itc}, will be modeled as changing exogenously as a function of the family's location, the child's grade, and the chosen action. The probability of being assigned to school m_{itc} is made up of two parts:

1. A new draw from the school assignment distribution, which depends on the child's grade and the family's location. Call this distribution $M_t^{New}(h_{it}, g_{itc})$
2. The probability of re-drawing from the school assignment distribution (I refer to this as the "re-assignment probability"). This probability is modeled as:

$$M^{RA}(d_{it}, g_{itc}, h_{i,t-1}) = \begin{cases} \rho & \text{if } d_{it} \in \{stay\} \cup \mathcal{R}\left(h_{i,t-1}\right) \text{ and } g_{itc} \notin \{1, 7, 9\} \\ 1 & \text{otherwise} \end{cases}$$

   where $\mathcal{R}\left(h_{i,t-1}\right)$ is the set of locations in the local region of $h_{i,t-1}$.

I model the school assignment in the following way. If family stays in their current location or moves to a nearby location (as captured by $\mathcal{R}\left(h_{i,t-1}\right)$), their child will stay at the same school as before with

[68] To be clear, the probabilities are assigned as a one third probability of being born between January and April, and two thirds probability of being born in May to December. The initial grades assigned at age a^* are negative purely for notational convenience, e.g. if $a^* = 0$, then a starting grade of -4 means that the child reaches grade 1 at age 5.

[69] In the school records, over all child-year observations, I find that children move up by exactly one grade in a subsequent year in 98.6% of observations.

probability $1-\rho$. The one exception is if the child begins a new school level, which occurs at grades 1, 7, and 9. In these cases – if the family moves far from their original home or if their child graduates to a new school level – then the child will receive a new school assignment draw with probability 1. If a new draw does occur, then it follows the distribution $M_t^{New}(h_{it}, g_{itc})$.

This means that the full school assignment probability is given by:

$$m_{itc} \sim M_t\left(d_{it}, g_{itc}, h_{i,t-1}\right) = \begin{cases} M_t^{New}(h_{it}, g_{itc}) & \text{with probability } M^{RA}(d_{it}, g_{itc}, h_{i,t-1}) \\ \mathbf{I}\left\{m_{i,t-1,c}\right\} & \text{otherwise} \end{cases} \tag{C.7}$$

where $\mathbf{I}\{m_{i,t-1,c}\}$ is the degenerate distribution for the previous school, $m_{i,t-1,c}$. The assignment distribution M_t^{New} is indexed by t as it is implicitly a function of the school zones and lottery process (which are all captured by ω_t). This function can therefore also capture changes over time such as the introduction of zones and the opening/closing of schools.

These modeling choices are supported by empirical patterns. In Figure C.2, I show that the probability of a school change among those who stay in their current home is low but non-zero (approximately 3-4% for most grades). This is evidence that $\rho > 0$.

Figure C.2: Probability of School Changing by Grade (Non-Movers)

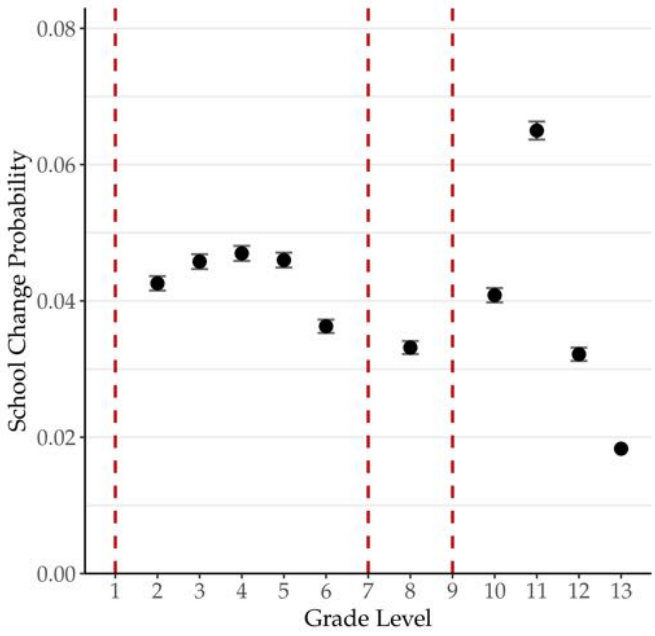

Note: The figure shows the proportion of child who attend a different school than the previous year, calculated by the grade of the child. This proportion is calculated only among families that stayed in the same home as the previous year. Grades 1, 7, and 9 are excluded as these are the starting grades for each school level (primary, intermediate, and high school, respectively).

Relatedly, Figure C.3 shows that the probability of a school change increases with distance, and that this is especially true for primary school students (likely because the catchment areas tend to be smaller). This suggests that moves to a nearby location should have a different likelihood of school re-assignment than moves to a distant location. Finally, note that the school assignment probability M_t does not differ by type ψ. Figure C.4 shows that the distribution of distances between home and school looks nearly identical across the three ethnic groups. This suggests that – conditional on location – their school attendance patterns are likely to be similar, which is evidence in support of this parameterization.

Figure C.3: Probability of School Changing by Move Distance

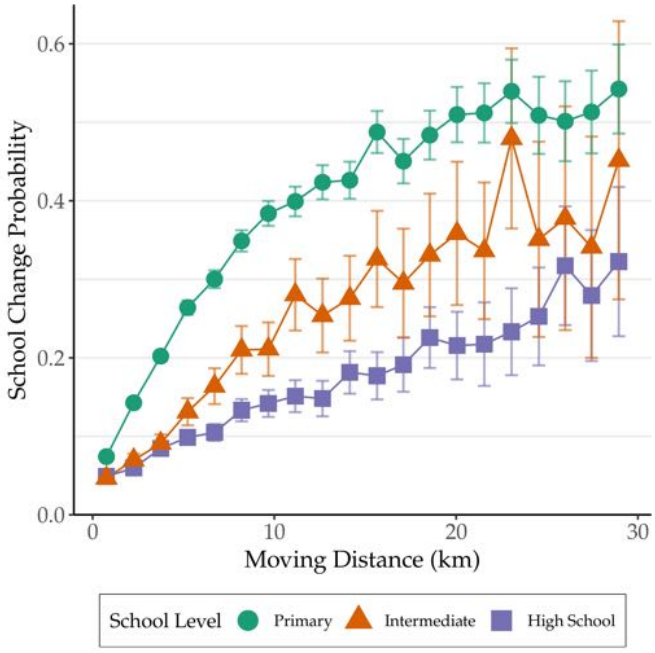

Note: The figure shows, among families who moved houses, the proportion of child who attend a different school than the previous year, relative to the distance of the move. The moving distance is calculated (in kilometers) as the straight line distance between Level 1 (Figure A.16a) centroids. This is calculated separately by school level, and only among students who continue to be in the same school level before and after the move.

As schools are a central part of the research question, it is worthwhile to discuss the caveats of modeling the school assignment process in this way. In the estimation, the school assignment distribution M_t^{New} will be estimated using the empirical school attendance probabilities conditional on where students live. This means that this distribution will capture a combination of the true admission probabilities (e.g. arising from school zones and lotteries) as well as parent preferences. In theory, these could be separately identified. However, doing so would require having additional data about the application process – for example, the set of schools that each family applied to and which of them they were admitted into.[70]

To see the potential issue with my modeling approach, consider the following example. Suppose I

[70]Due to the decentralized nature of the school system, such data is not collected by the New Zealand Ministry of Education.

Figure C.4: Distance to School Distribution, by Ethnicity

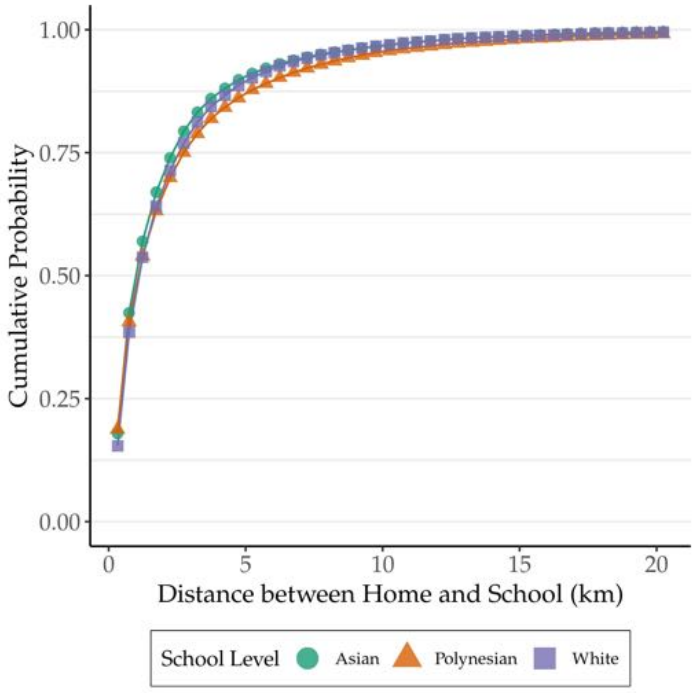

Note: The figure shows the cumulative distribution of the distance between home and school. Distance is calculated in kilometers as the straight line distance between the Level 1 (Figure A.16a) centroid of home location and the school coordinates. This is calculated separately by ethnic group.

observe all families living in location A sending their children to school Z. Would it be reasonable to assume that *any* family who then moves to location A will also send their child to school Z? This is a plausible assumption if non-individual factors such as admissions zones or distance largely determine the choice of school. If it is instead driven by some other parent preferences, then making such an extrapolation may bias the estimates. This is because families who move to location A in a counterfactual world may not share the same preferences as those who we observe in the data as living in location A.

Figure C.5 shows that most zoned schools have a high proportion of their students living inside their zone, and that this is especially true at the highest and lowest deciles. Figure A.10 also shows that children tend to attend the schools that are closest to their homes, emphasizing the importance of distance. As these factors (zoning and distance) appear to play an important role in the school decision process, it seems that location and observed attendance patterns can reasonably approximate the school choice process.

Figure C.5: Fraction of Students Living In-Zone

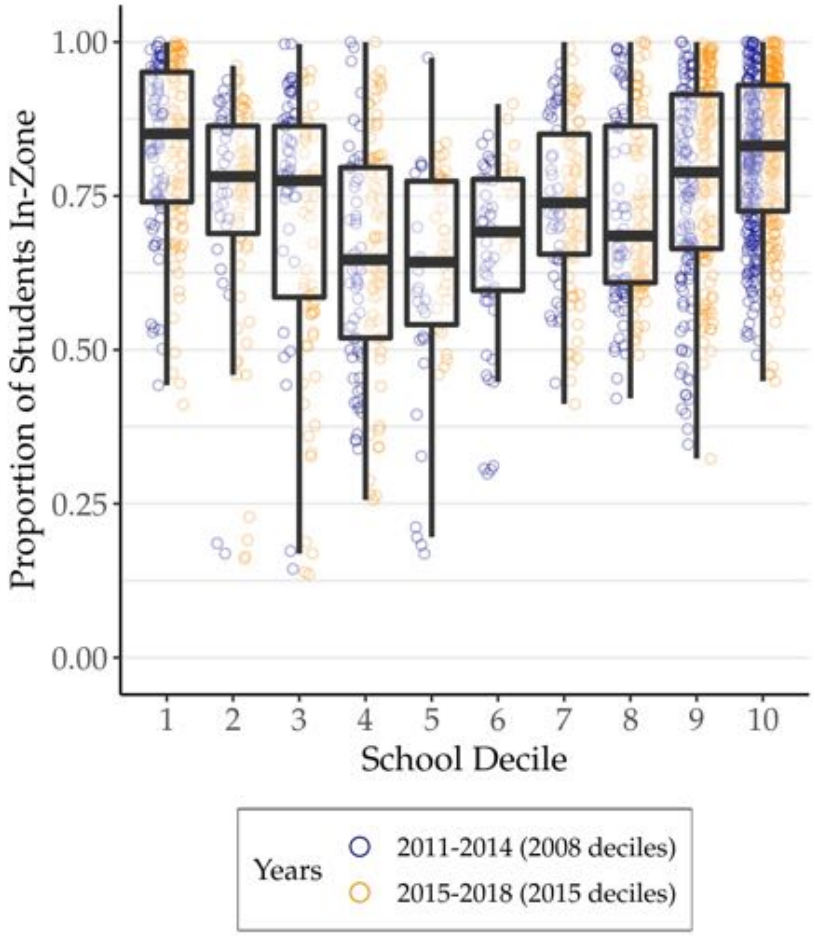

Note: The figure shows the the proportion of student living inside a school's zone (vertical axis) relative to the decile score of the school (horizontal axis). The decile score values receive a small perturbation to show variation. Each observation represents a school-year for the period 2011 to 2018. The 2011-2014 observations (blue) have decile scores there were assigned in 2008. The 2015-2018 observations (orange) have decile scores there were assigned in 2015. Box plots are overlaid for each decile score to show the 25th percentile, median, and 75th percentile across all school-years within the decile score.

C.2 Supporting Figures

Figure C.6: Spatial Distribution of Population Subgroups (2018)

(a) Parents

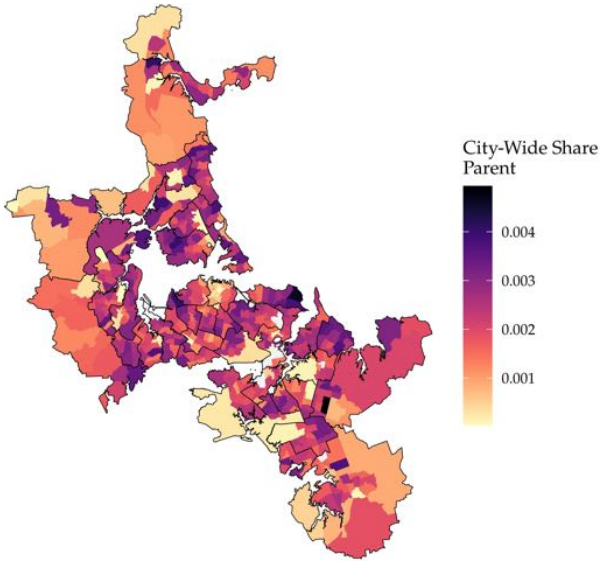

(b) No Dependents, Aged Under 45

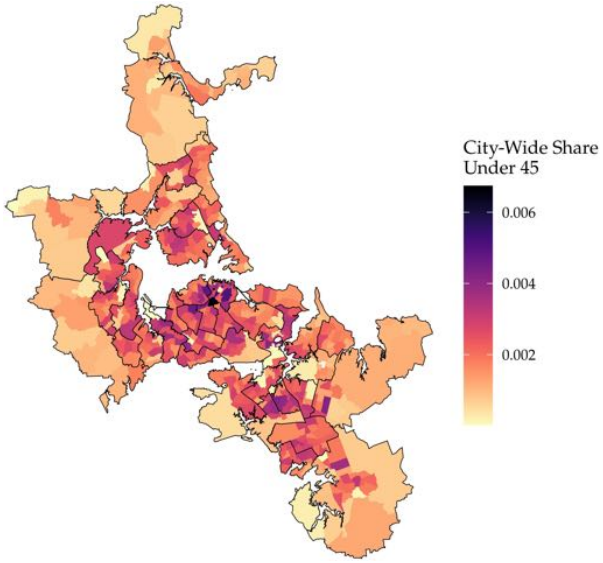

(c) No Dependents, Aged 45-65

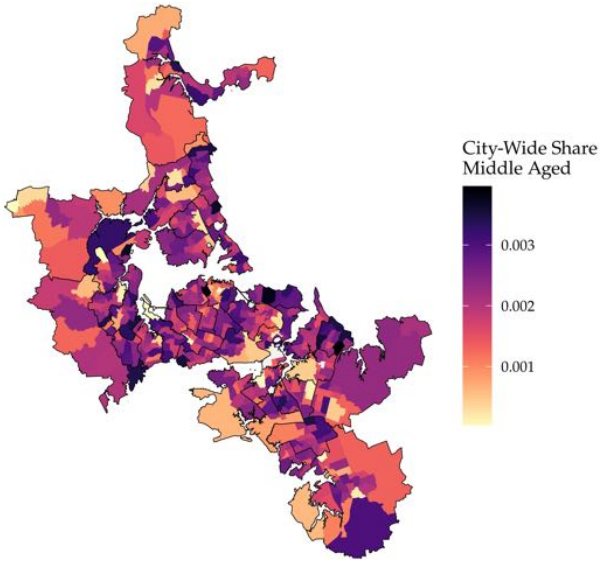

(d) No Dependents, Aged Over 65

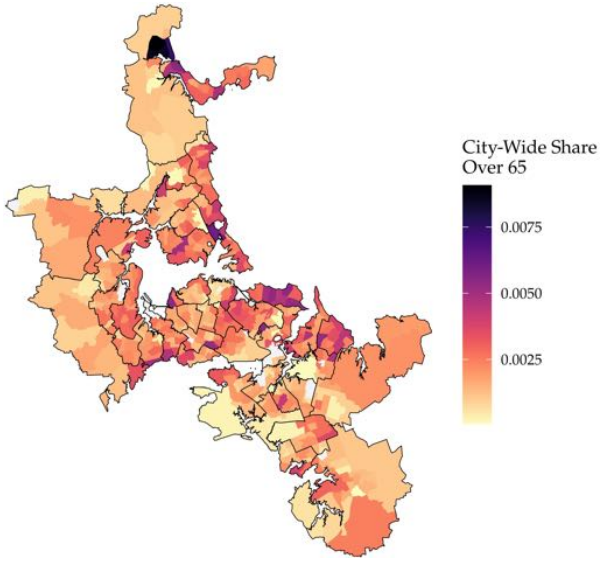

Note: The maps show the spatial distribution of groups within the Auckland urban area using the 2018 census. For each group, I shade each location (Level 2 unit, Figure A.16b) according to the fraction of the group's total (city-wide) population that is living in that location. The groups are parents (those living with dependent children) and three groups of non-parents (over 18 year-olds without dependents). The non-parents are grouped by their age.

Figure C.7: Moving Probability versus Tenure

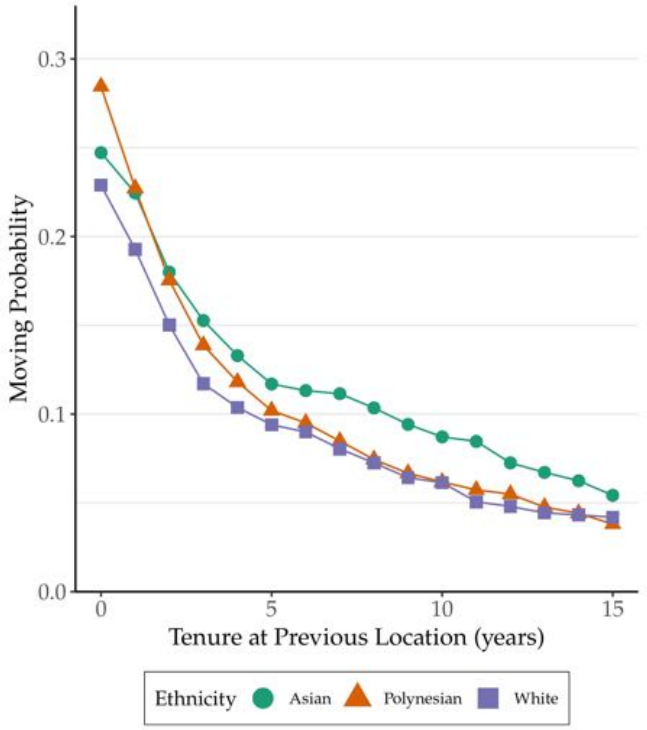

Note: The figure shows the mean proportion of families moving given the number of years lived in their current home. This is calculated separately by ethnic group.

Figure C.8: Moving Distance Distribution

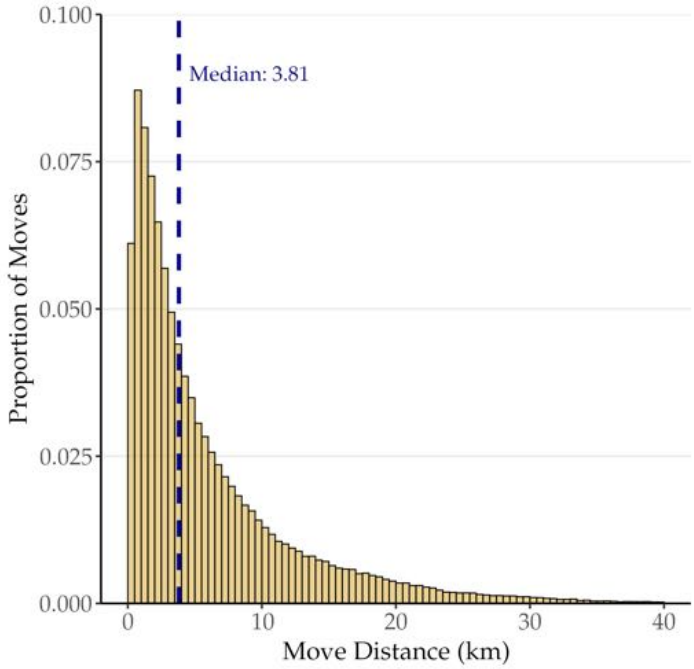

Note: The figure shows the distance between the previous and new home location among households who move within Auckland city. The moving distance is calculated as the straight line distance between Level 1 (Figure A.16a) centroids. Each bin represents 2 kilometers.

Figure C.9: Moving Distance versus Tenure

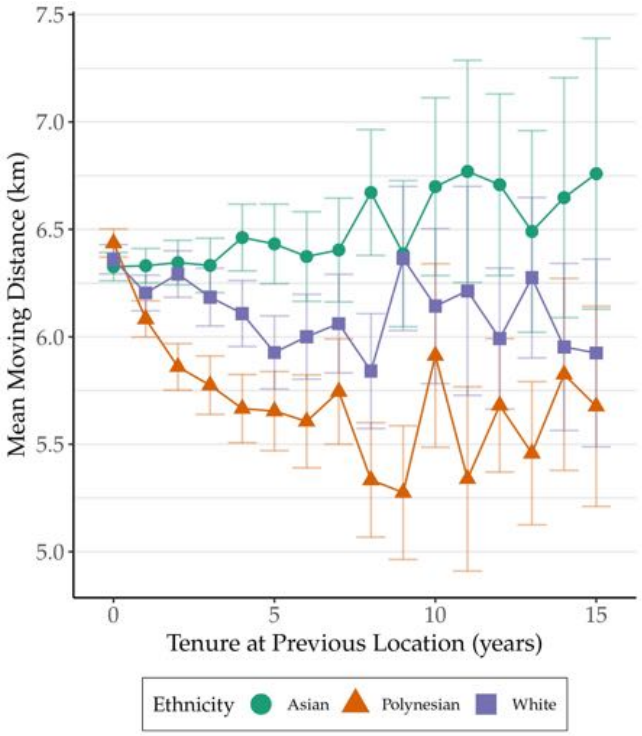

Note: The figure shows, conditional on moving, the mean moving distance (in kilometers) given the number of years lived in previous home. The moving distance is calculated as the straight line distance between Level 1 (Figure A.16a) centroids. This is calculated separately by ethnic group.

C.3 ECCP Illustration

Consider a family living in location O at the beginning of period t. For period $t+1$, suppose they can choose (d_t) between moving to location A or to location B, which is chosen depending on the value received from each action. The empirical challenge in estimating a dynamic problem is in separating the flow utility from the continuation value for each action's payoff.

The ECCP methodology overcomes this challenge. Now, regardless of their chosen d_t action, suppose that the family chooses (d_{t+1}) to move to location C in period $t+2$. We can then compare the payoff of the action tuple $(d_t, d_{t+1}) = (A, C)$ to that of (B, C), as illustrated in Figure C.10. The key insight is that upon arriving in location C, the family's continuation value is the same regardless of their past actions. In either case, the family is new to location C, which means they will lose any of their previously accumulated tenure capital and would have to re-enroll their children in new schools.

Figure C.10: ECCP Example Diagram

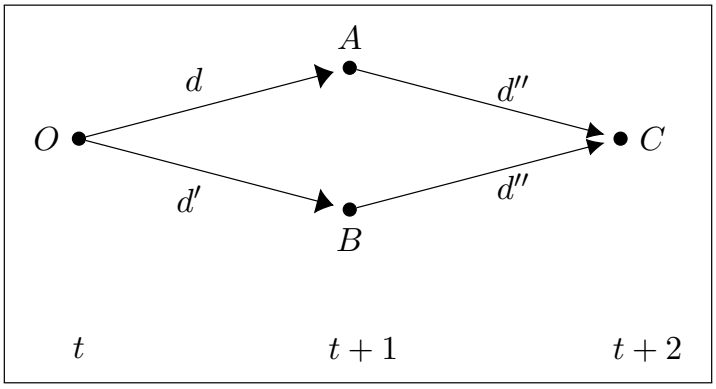

Note: The diagram illustrates the ECCP methodology for this setting. In period t, the family lives in location O. One option is to choose action $d_t = d$ that leads to living in location A in period $t+1$. Another option is $d_t = d'$, which leads to living in location B in $t+1$. Following this, an action $d_{t+1} = d''$ in period $t+1$ that leads to living in location C in period $t+2$ means that the family will end up in the same state regardless of their past choice in period t.

As the value of C at $t+2$ is identical regardless of their past behavior, that means that if we compare the payoff of the action path $O \to A \to C$ versus the path $O \to B \to C$, then the *difference* will only be in terms of payoffs that occur between O and C. All future expected payoffs from $t+2$ onwards will get cancelled out when taking this difference. With this, we can then conclude that the relative likelihood of observing the action path $(O \to A \to C)$ over $(O \to B \to C)$ can be explained by three components:

Relative likelihood of $(O \to A \to C)$ versus $(O \to B \to C)$	=	Difference in moving cost of $(O \to A)$ versus $(O \to B)$	+	Difference in flow utility of living in A versus B	+	Difference in moving cost of $(A \to C)$ versus $(B \to C)$

C.4 ECCP Derivation

This derivation follows sections 2.1 and 3 in Kalouptsidi et al. (2020). Variables and notation definition follow those described in section 4.

C.4.1 Setup

The value function of the family's maximization problem at time t is:

$$V(\mathbf{s}_{it}) = \max_{d \in \mathcal{D}} \{\Pi(d, \mathbf{s}_{it}) + \beta E_{\mathbf{s}} \left[V(\mathbf{s}_{i,t+1}) | \, d, \mathbf{s}_{it} \right]\}$$

where $\Pi(d, \mathbf{s}_{it})$ is the per period utility defined in Equation (3).

The ex-ante value function is defined as the expectation of the value function, taken over the idiosyncratic errors ε_{it}:

$$\tilde{V}(k_{it}, \omega_t) \equiv E_{\varepsilon} \left[V(k_{it}, \omega_t, \varepsilon_{it}) \right]$$

The choice-specific (or conditional) value function of an action d is defined as:

$$v_d(k_{it}, \omega_t) \equiv u(d, k_{it}, \omega_t) + \beta E \left[\tilde{V}(k_{i,t+1}, \omega_{t+1}) \Big| \, d, k_{it}, \omega_t \right] \tag{C.8}$$

With Type 1 extreme value errors (Assumption 3), the conditional choice probability of an action d is given by:

$$\begin{aligned} p_d(k_{it}, \omega_t) &= Pr \left(d_{it} = d | \, k_{it}, \omega_t \right) \\ &= \frac{\exp \left(v_d(k_{it}, \omega_t) \right)}{\sum_{d' \in \mathcal{D}} \exp \left(v_{d'}(k_{it}, \omega_t) \right)} \end{aligned} \tag{C.9}$$

And the ex-ante value function can be expressed as:

$$\tilde{V}(k_{it}, \omega_t) = \log \left(\sum_{d' \in \mathcal{D}} \exp v_{d'}(k_{it}, \omega_t) \right) + \gamma \tag{C.10}$$

where γ is Euler's constant.

For a fixed action d, (C.9) can be substituted into (C.10) to express the ex-ante value function as:

$$\tilde{V}(k_{it}, \omega_t) = v_d(k_{it}, \omega_t) - \ln p_d(k_{it}, \omega_t) + \gamma \tag{C.11}$$

Note that (C.11) holds for any action $d \in \mathcal{D}$.

Lastly, the expectational error is defined. For a true realization of the market state variables $\omega_{t+1} \in \Omega$, the expectational error e_d^V is given by:

$$\begin{aligned} e_d^V(k_{it}, \omega_t) &= \sum_{k' \in \mathcal{K}} \underbrace{\left[E_{\omega'} \left[\tilde{V}(k', \omega') \middle| \omega_t \right] - \tilde{V}(k', \omega_{t+1}) \right]}_{\substack{\text{prediction error for a} \\ \text{realized value of } k'}} F^k \left(k' \middle| d, k_{it}, w_t \right) \\ &= E_{k,\omega} \left[\tilde{V}(k_{i,t+1}, \omega_{t+1}) \middle| d, k_{it}, \omega_t \right] - E_k \left[\tilde{V}(k_{i,t+1}, \omega_{t+1}) \middle| d, k_{it} \right] \end{aligned} \tag{C.12}$$

where $k' \in \mathcal{K}$ and $\omega' \in \Omega$ denote the possible next period $(t+1)$ values for k_t and w_t. The expectation error is calculated as the corresponding prediction error in ω for a given k', and then integrating this over the realizations k' of the state variable k_{t+1}. In the second line, notice that the first term is an expectation over k and ω, while the second one is only over k.

C.4.2 Flow Utility

Using the definition of choice-specific value function (Equation C.8), we can express the flow utility as follows:

$$u_{dt}(k_{it}) = v_{dt}(k_{it}) - \beta E_{k,\omega} \left[\tilde{V}_{t+1}(k_{i,t+1}) \middle| d, k_{it}, \omega_t \right] \tag{C.13}$$

where the d and t subscripts are used as a shorthand to eliminate writing d and ω_t as an input to a function (similarly, a $t+1$ indicates ω_{t+1} as an input).

The first term of (C.13), the choice specific value function, can be expressed using Equation (C.11):

$$v_{dt}(k_{it}) = \tilde{V}_t(k_{it}) - \gamma + \ln p_{dt}(k_{it}) \tag{C.14}$$

The second term of (C.13), the expected continuation value, is given by re-arranging Equation (C.12):

$$\begin{aligned} E_{k,\omega} \left[\tilde{V}_{t+1}(k_{i,t+1}) \middle| d, k_{it}, \omega_t \right] &= E_k \left[\tilde{V}_{t+1}(k_{i,t+1}) \middle| d, k_{it} \right] + e_{dt}^V(k_{it}) \\ &= \sum_{k' \in \mathcal{K}} \tilde{V}_{t+1}(k') F^k \left(k' \middle| d, k_{it}, w_t \right) + e_{dt}^V(k_{it}) \\ &= F_{dt}^k(k_{it}) \tilde{\mathbf{V}}_{t+1} + e_{dt}^V(k_{it}) \end{aligned} \tag{C.15}$$

where $\tilde{\mathbf{V}}_{t+1} = \left[\tilde{V}_{t+1}(k) \right]_{k \in \mathcal{K}}$ is a $K \times 1$ vector and $F_{dt}^k(k_{it})$ is a $1 \times K$ transition vector defined in section 4.5.1.[71]

[71] In other words, $F_{dt}^k(k_{it} = k^*) = \left[F^k \left(k_{i,t+1} = k_1 \middle| d, k_{it} = k^*, w_t \right), \cdots, F^k \left(k_{i,t+1} = k_K \middle| d, k_{it} = k^*, w_t \right) \right]$

Substituting (C.14) and (C.15) into (C.13) gives:

$$u_{dt}(k_{it}) = \tilde{V}_t(k_{it}) - \gamma + \ln p_{dt}(k_{it}) - \beta F^k_{dt}(k_{it})\tilde{\mathbf{V}}_{t+1} - \beta e^V_{dt}(k_{it})$$

To further simplify the notation, I omit the k_{it} input for brevity. Throughout the remainder of the derivation, we should consider k_{it} as fixed to some value k^* (as in section 4.5.1). This means the flow utility is expressed as:

$$u_{dt} = \tilde{V}_t - \gamma + \ln p_{dt} - \beta F^k_{dt}\tilde{\mathbf{V}}_{t+1} - \beta e^V_{dt} \tag{C.16}$$

C.4.3 ECCP Equations

Starting from the initial point $k_{it} = k^*$, using Equation (C.16), consider two different actions d and d':

$$\begin{aligned} u_{dt} &= \tilde{V}_t - \gamma + \ln p_{dt} - \beta F^k_{dt}\tilde{\mathbf{V}}_{t+1} - \beta e^V_{dt} \\ u_{d't} &= \tilde{V}_t - \gamma + \ln p_{d't} - \beta F^k_{d't}\tilde{\mathbf{V}}_{t+1} - \beta e^V_{d't} \end{aligned}$$

Taking the difference of the above equations and re-arranging:

$$\begin{aligned} u_{dt} - u_{d't} =& \ln p_{dt} - \ln p_{d't} - \beta\left(F^k_{dt} - F^k_{d't}\right)\tilde{\mathbf{V}}_{t+1} - \beta\left(e^V_{dt} - e^V_{d't}\right) \\ \ln\frac{p_{dt}}{p_{d't}} =& u_{dt} - u_{d't} + \beta\left(F^k_{dt} - F^k_{d't}\right)\tilde{\mathbf{V}}_{t+1} + \beta\left(e^V_{dt} - e^V_{d't}\right) \end{aligned} \tag{C.17}$$

Next, we find an action d'' that is a renewal action, as defined in section 4.5.1. Formally, a renewal action d'' (taken at $t+1$) satisfies the following property given the two actions d and d' (taken at t):[72]

$$F^k_{dt}\mathbf{F}^k_{d'',t+1} = F^k_{d't}\mathbf{F}^k_{d'',t+1} \tag{C.18}$$

where $\mathbf{F}^k_{dt} = \left[F^k_{dt}(k)\right]_{k\in\mathcal{K}}$ is the $K \times K$ transition matrix defined in section 4.5.1. Recall that $F^k_{dt} = F^k_{dt}(k^*)$ as we are implicitly keeping k_{it} fixed at k^*.

Consider Equation (C.16) iterated forward by one period to $t+1$ and with the chosen action as $d_{i,t+1} = d''$:

$$u_{d'',t+1} = \tilde{V}_{t+1} - \gamma + \ln p_{d'',t+1} - \beta F^k_{d'',t+1}E_{\omega|t+1}\left[\tilde{\mathbf{V}}_{t+2}\right] \tag{C.19}$$

The continuation value expressed in (C.19) is simply an alternative formulation of the iterated-forward

[72] Kalouptsidi et al. (2020) define a renewal action as an action that if taken "in period $t+1$ leads to the same distribution of states at the beginning of time period $t+2$, *regardless of which state the agent was in during period* t" (emphasis mine). This definition does not exactly apply in this situation, given that renewal actions are chosen based on state k_{it} and past actions (d, d') to ensure Equation (C.18) holds. However, I will use the term here as it conveys the same intuition as in other cases where renewal actions are used (Almagro and Domínguez-Iino, 2021; Scott, 2013; Rust, 1987).

version of (C.15).[73]

To simplify notation, define the following $K \times 1$ vectors: $\mathbf{u}_{d'',t+1} = \left[u_{d'',t+1}(k)\right]_{k\in\mathcal{K}}$, $\boldsymbol{\gamma} = \left[\gamma\right]_{k\in\mathcal{K}}$ and $\ln\mathbf{p}_{d'',t+1} = \left[\ln p_{d'',t+1}(k)\right]_{k\in\mathcal{K}}$. Since $\tilde{\mathbf{V}}_{t+1}$ is a vector of $\tilde{V}_{t+1}(k)$ terms, this means we can use Equation (C.19) to express it as:

$$\begin{aligned}\tilde{\mathbf{V}}_{t+1} &= \left[u_{d'',t+1}(k) + \gamma - \ln p_{d'',t+1}(k) + \beta F^k_{d'',t+1}(k) E_{\omega|t+1}\left[\tilde{\mathbf{V}}_{t+2}\right]\right]_{k\in\mathcal{K}} \\ &= \mathbf{u}_{d'',t+1} + \boldsymbol{\gamma} - \ln\mathbf{p}_{d'',t+1} + \beta \mathbf{F}^k_{d'',t+1} E_{\omega|t+1}\left[\tilde{\mathbf{V}}_{t+2}\right]\end{aligned}$$

Recalling that k_{it} is fixed at k^*, this means we can derive the following relationship:

$$\left(F^k_{dt} - F^k_{d't}\right)\tilde{\mathbf{V}}_{t+1} = \left(F^k_{dt} - F^k_{d't}\right)\left(\mathbf{u}_{d'',t+1} - \ln\mathbf{p}_{d'',t+1}\right) \tag{C.20}$$

This is because by the renewal action property, we have:[74]

$$F^k_{dt}\mathbf{F}^k_{d'',t+1} E_{\omega|t+1}\left[\tilde{\mathbf{V}}_{t+2}\right] = F^k_{d't}\mathbf{F}^k_{d'',t+1} E_{\omega|t+1}\left[\tilde{\mathbf{V}}_{t+2}\right]$$

Plugging Equation (C.20) into Equation (C.17) gives us:

$$\begin{aligned}\ln\frac{p_{dt}}{p_{d't}} &= u_{dt} - u_{d't} + \beta\left(F^k_{dt} - F^k_{d't}\right)\left(\mathbf{u}_{d'',t+1} - \ln\mathbf{p}_{d'',t+1}\right) + \beta\left(e^V_{dt} - e^V_{d't}\right) \\ \ln\frac{p_{dt}}{p_{d't}} + \beta\left(F^k_{dt} - F^k_{d't}\right)\ln\mathbf{p}_{d'',t+1} &= \bar{u}_{dt} + \xi_{dt} - \bar{u}_{d't} - \xi_{d't} + \beta\left(F^k_{dt} - F^k_{d't}\right)\mathbf{u}_{d'',t+1} + \beta\left(e^V_{dt} - e^V_{d't}\right)\end{aligned}$$

Therefore, the regression to estimate is:

$$\ln\frac{p_{dt}}{p_{d't}} + \beta\left(F^k_{dt} - F^k_{d't}\right)\ln\mathbf{p}_{d'',t+1} = \bar{u}_{dt} - \bar{u}_{d't} + \beta\left(F^k_{dt} - F^k_{d't}\right)\mathbf{u}_{d'',t+1} + \epsilon_{dd't} \tag{C.21}$$

where the econometric error term is defined as:

$$\epsilon_{dd't} = \tilde{\xi}_{dd't} + \tilde{e}^V_{dd't} \quad \text{with} \quad \tilde{\xi}_{dd't} = \xi_{dt} - \xi_{d't} \quad \text{and} \quad \tilde{e}^V_{dd't} = \beta\left(e^V_{dt} - e^V_{d't}\right)$$

Since d'' is a renewal action, then $u_{d'',t}(k) - SMC_{dt}(k) - MC_{d''}(k)$ is equal for all $k \in \mathcal{K}$ (utility is the same excluding the moving costs). Call this value $\tilde{u}$. This means that the following holds:[75]

$$\begin{aligned}\left(F^k_{dt} - F^k_{d't}\right)\mathbf{u}_{d'',t+1} &= \left(F^k_{dt} - F^k_{d't}\right)\left[\tilde{\mathfrak{u}} + SMC_{dt}(k) + MC_{d''}(k)\right]_{k\in\mathcal{K}} \\ &= SMC_{d,d'',t} - SMC_{d,d'',t} + MC_{d,d''} - MC_{d',d''}\end{aligned}$$

[73] Specifically, $E_{k,\omega}\left[\tilde{V}_{t+2}(k_{i,t+2})\middle| d'', k_{i,t+1}, \omega_{t+1}\right] = F^k_{d'',t+1}\left[E_\omega\left[\tilde{V}_{t+2}(k)\middle|\omega_{t+1}\right]\right]_{k\in\mathcal{K}} = F^k_{d'',t+1}E_{\omega|t+1}\left[\tilde{\mathbf{V}}_{t+2}\right]$

[74] The γ term also goes away since the following holds for any d: $F^k_{dt}\boldsymbol{\gamma} = \gamma$.

[75] Footnote 74 gives a similar logic for this step. Footnote 25 defines the terms in the second line.

D Estimation Appendix

D.1 State Space Choices

To make the state space manageable, I restrict to families who have at most two children: $C_i \in \{1, 2\}$. I also restrict the maximum gap between the first and second child to be 10 years. I set the age of grade assignment as $a^* = 4$. This model component is discussed in Appendix C.1.4. The parameter choice is also motivated by the data. In Figure B.16, I observe the moving probabilities for the two grade cohorts diverging when they are both age 4. After age 4, it is the grade of the child (and not their age) that affects the families choice probabilities.

I set $\underline{b} = 3$ and $\bar{b} = 25$, which means I look at the period from 3 years prior to the eldest child's birth to when the youngest child reaches 25. I set $\underline{b} = 3$ as I observe a jump in moving probability at this point in Figure B.16. As the average mother's age at the time of a child's birth is 35, I set $\bar{b} = 25$ as this corresponds to when the average mother reaches age 60. While I use data on each family until their youngest child turns 25. However, since children past high school are less relevant for my model (and are fewer in numbers), I additionally simplify the state space by capping age to 22 and grade to 14. All children with values above these are pooled together for the purposes of estimation.

Given these restrictions, I observe 674 possible child-age-grade combinations.

To set the locations, I need to discretize the city area. I combine geographical units defined by Statistics New Zealand to create a hierarchy of distinct geographic areas as shown in Figure A.16. Level 1 and 2 units are defined by Statistics New Zealand (referred to as meshblocks and statistical area 2 (SA2) units, respectively). The remainder are my own creation based on a grouping algorithm that takes into account population size and local community board boundaries. The key point is that there is a many-to-one match from each Level A unit to the subsequent Level $A+1$ unit, i.e. each geographic level is constructed using the previous level's units.

Locations are set as the Level 3 areas (Figure A.16c). There are $|\mathcal{L}| = 42$ locations in the city. In 2018, the population of these locations ranged from 15,867 to 65,130 people, with a median population of 32,264. Regions $\mathcal{R}(\cdot)$ are set as Level 5 areas (Figure A.16e), where there are 14 in the city. Each region has approximately $|\mathcal{R}(h_i)| = 3$ locations. The outside option is set as any location outside the Auckland urban area, including overseas.[76] This means that there are $|\mathcal{D}| = 2 + |\mathcal{L}| = 44$ possible actions (each location plus the stay and outside option actions).

[76]For people who first immigrate to New Zealand, I do not include the years prior to their immigration as part of the estimation.

D.2 CCP Estimation Details

D.2.1 CCP Grouping Algorithm

In this section, I describe how I applied the grouping algorithm of Raval et al. (2017) to this setting.

Each observation corresponds to a state $\mathbf{s}_{it}$. I map each state into a set of features. These are partitioned into two groups: immutable and mutable features.

$$f(\mathbf{s}_{it}) = \Big(\underbrace{f_1 \quad \cdots \quad f_N}_{\substack{\text{Immutable} \\ \text{features}}} \Big| \underbrace{f_1 \quad \cdots \quad f_M}_{\substack{\text{Mutable} \\ \text{features}}} \Big)$$

The algorithm aims to group observations with similar features. The requirement is that groups must meet a minimum group size threshold, which I will call S_{min}. Observations are grouped in an iterative process. If a group of observations meets the minimum group size, it is designated as a final group and removed from the algorithm process. For all remaining observations, I drop their lowest ranked mutable feature and repeat the process until the mutable features are exhausted. In other words, each iteration, observations that are not part of large enough groups lose one level of specificity and are re-pooled with otherwise similar observations in the hopes of making a large enough group. This process is summarized in Algorithm 1.

Algorithm 1 Grouping Algorithm

Assign each observation their set of features $f(\mathbf{s}_{it})$, with N immutable features and M mutable features
for $m = 0$ to M **do**
Group remaining observations based on the N immutable features and the first $M - m$ mutable features, i.e. $f_1, ..., f_{M-m}$. For $m == M$", use no mutatable features.
Allocate group identifiers to groups that have a size S_g that meet the minimum threshold, i.e. $S_g \geq S_{min}$
Remove observations with a group identifier
end for
Drop any remaining observations without a group identifier

To implement the above algorithm, I need to set the features and the minimum group size. I use a threshold of $S_{min} = 25$. To allow for more specific groupings, I bin several variables to increase the number of features to choose from. The feature order I choose, from highest to lowest ranked, is as follows:

Grouping Algorithm Features

Immutable Features $(1-N)$

1. Non-binned variables:
   (a) Type $\psi \in \Psi$
   (b) Number of dependents $\left(\sum_{c=1}^{C_{it}} \mathbb{1}\{a_{ict} \leq 22\}\right) \in \{0,1,2\}$
2. Broadly binned variables:
   (a) Broadly binned age/grade:†
   - $a_{B2} \in \{\{-3\},[-2,0],[1,3],[4,\infty)\}$
   - $g_{B2} \in \{[-8,-2],[-1,1],[2,5],[6,7],[8,9],[10,13],[14,\infty)\}$

   (b) Broadly binned year: $t_{B2} \in \{[2008,2013],[2014,2019]\}$
   (c) Broadly binned home: level 4 locations, see Figure A.16d

Mutable Features $(1-M)$

1. Binned tenure:* $\tau_B \in \{[0,1],[2,3],[4,5]\}$
2. Home: $h \in \mathcal{L} \cup \{0\}$, level 3 locations, see Figure A.16c
3. Binned year: $t_B \in \{[2008,2010],[2011,2013],[2014,2016],[2017,2019]\}$
4. Binned age/grade of the child:†
   - $a_B \in \{\{-3\},[-2,-1],\{0\},[1,2],\{3\},[5,\infty)\}$
   - $g_B \in \{[-8,-2],\{-1\},\{0\},\{1\},[2,5],\{6\},\{7\},[8,9],[10,11],[12,13],[14,\infty)\}$
5. Year: $t \in \{2008,...,2019\}$
6. Age/grade of the child:† $a \in \{-3,...,21,[22,\infty)\}$, $g \in \{-8,...,13,[14,\infty)\}$
7. Tenure: $\tau \in \{0,1,2,3,4,5\}$

*For the first step CCP estimation $\left(p_t^{SMO}\right)$, I set the binned tenure τ_B as an immutable feature.
†Age is the feature until age $a^* = 4$, then grade becomes the feature after that. This is chosen to be consistent with the estimation. Age is capped at $a = 22$. Grade is capped at $g = 14$ for all those who finish school.

Table D.1 shows how the observations are allocated across the groups

Table D.1: Grouping Algorithm Outcome

		Number of Grouped Observations					
		First Step CCP			Second Step CCP		
Iteration	**Last Dropped Feature**	White	Polynesian	Asian	White	Polynesian	Asian
0	None	445,449	273,813	161,715	15,027	4,032	27,438
1	Tenure	52,617	32,931	44,091	3,951	2,166	3,234
2	Age/Grade	219,711	122,916	115,359	6,666	3,942	7,407
3	Year	318,117	226,809	254,169	26,595	22,494	23,838
4	Age/Grade (binned)	96,804	95,649	98,322	12,837	12,939	10,497
5	Year (binned)	65,823	68,457	63,249	26,388	20,277	21,342
6	Home	26,673	32,127	29,766	18,747	15,933	16,761
7	Tenure (binned)	–	–	–	28,272	27,945	30,156
Summary	Grouped Observations	1,225,194	852,702	766,671	138,483	109,728	140,673
	Ungrouped Observations	147,066	172,272	124,476	30,867	35,826	28,254
	Total Observations	1,372,260	1,024,974	891,147	169,350	145,554	168,927
	Percent Grouped (%)	89.3	83.2	86.0	81.8	75.4	83.3

Note: The table shows the outcome of the grouping algorithm. Each row in the first panel indicates the number of observations that were grouped in each iteration. An observation is a family-year observation from the family panel data. An iteration m is defined by having $M - m$ mutable features, i.e. the $(M - m + 1)^{th}$ feature is the "last dropped". Columns represent the CCP step and ethnic group, as each algorithm is applied separately for each of these. In the first step CCP, binned tenure is treated as an immutable feature. In the second step CCP, observations are restricted to those where the family moves. Total observations is equal to grouped plus ungrouped observations. Percent grouped is equal to grouped observations divided by total observations. In accordance with Statistics New Zealand confidentiality rules, the number of observations have been randomly rounded to a multiple of 3 (RR3).

D.2.2 CCP Two-Step Estimation

I estimate the CCPs using two steps. Within each of these steps, I apply the grouping algorithm described in D.2.1. The two steps are:

1. Estimate the conditional choice probability of observing the actions $d \in \{stay, move, outside\}$, where $move$ indicates moving to any location in the city, $l \in \mathcal{L}$. This differs to $outside$, which is moving to the outside option. Call this probability $p_t^{SMO}(d)$.
2. Among those who move to a location with the city, estimate the conditional choice probability of moving to each location $l \in \mathcal{L}$. Call this probability $p_t^M(d)$ for $d \in \mathcal{L}$.

With this, the CCPs can be expressed as:[77]

$$p_{dt} = \begin{cases} p_t^{SMO}(stay) & \text{if } d = stay \\ p_t^{SMO}(move)p_t^M(d) & \text{if } d \in \mathcal{L} \\ p_t^{SMO}(outside) & \text{if } d = 0 \end{cases}$$

For the first step CCPs (p_t^{SMO}), I calculate the CCP as the relative frequency of each action within each grouping of states. The scope for bias is much more limited in this step as the number of actions is reduced to three. Figure D.1 shows the distribution for the p_t^{SMO} CCPs across all the estimated states. As expected, the most likely action across states is $stay$, with a wide range of probabilities for $move$. The outside option tends to be a low probability action.

For the second step CCPs (p_t^M), I now leverage that the action set can be represented as spatial data (since this step is conditional on moving to a location inside the city). This helps overcome the dual issues of a smaller number of observations (I am restricting to in-city moves) as well as a large number of actions ($|\mathcal{L}| = 32$). Within each group, I use the coordinates of the new home location to conduct a spatial kernel density estimation (KDE). This uses the intuition that we should expect areas that are next to each other to have similar probabilities of being moved into. Figure D.2 illustrates two examples of the spatial KDE. To calculate the CCP for a location l, I evaluate the estimated spatial density function at each Level 1 centroid (Figure A.16a), then sum over all centroids within location l.

In Figure D.3, I compare the estimated CCPs to the move distance (i.e. the probability of choosing a move location given how far the family would be moving from their current home). This gives the expected shape wherein moves are (on average) more likely if they are closer to the family's home location. Despite the high dimensional problem and the simplifications made in the estimation process, Figures D.1 and D.3 reassuringly show that the CCPs seem to be reasonably reflecting the data.

[77] Recall that these values are holding fixed $k_{it} = k^*$. Therefore, $p_t^{SMO}(d)$ and $p_t^M(d)$ are each scalars.

Figure D.1: Stay-Move-Outside CCP

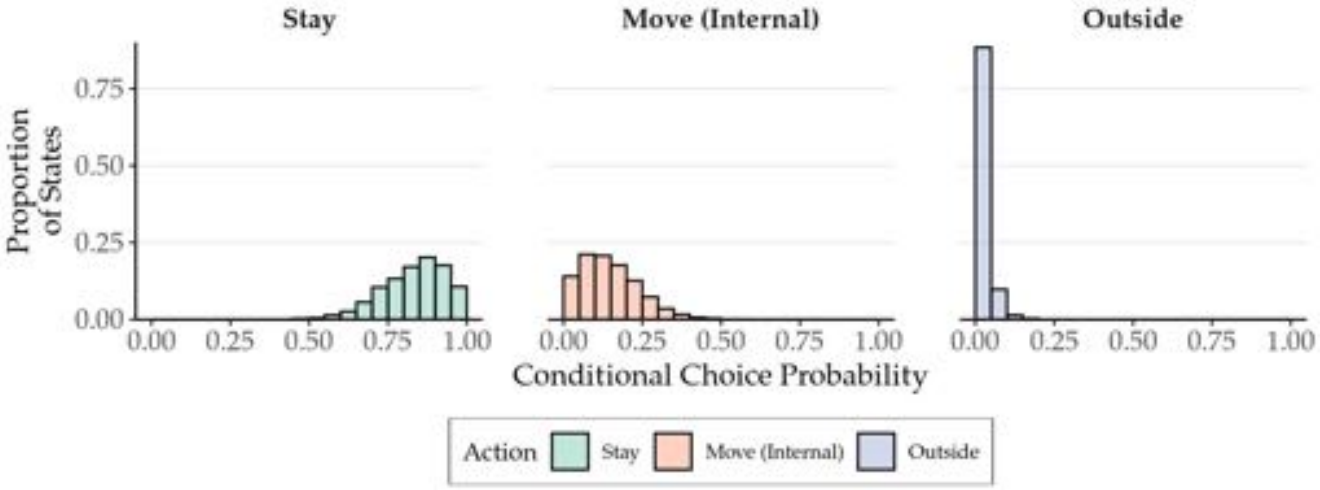

Note: The figures show the distribution of first step conditional choice probabilities (CCPs), where the choices are stay, move (to a location inside the city), or move to the outside option. The distribution is taken over the number of states that fall into each CCP bin. Bins are in 5 percentage point increments.

Figure D.2: Spatial KDE for Internal Move CCPs

(a) Example 1 (internal home) **(b)** Example 2 (outside home)

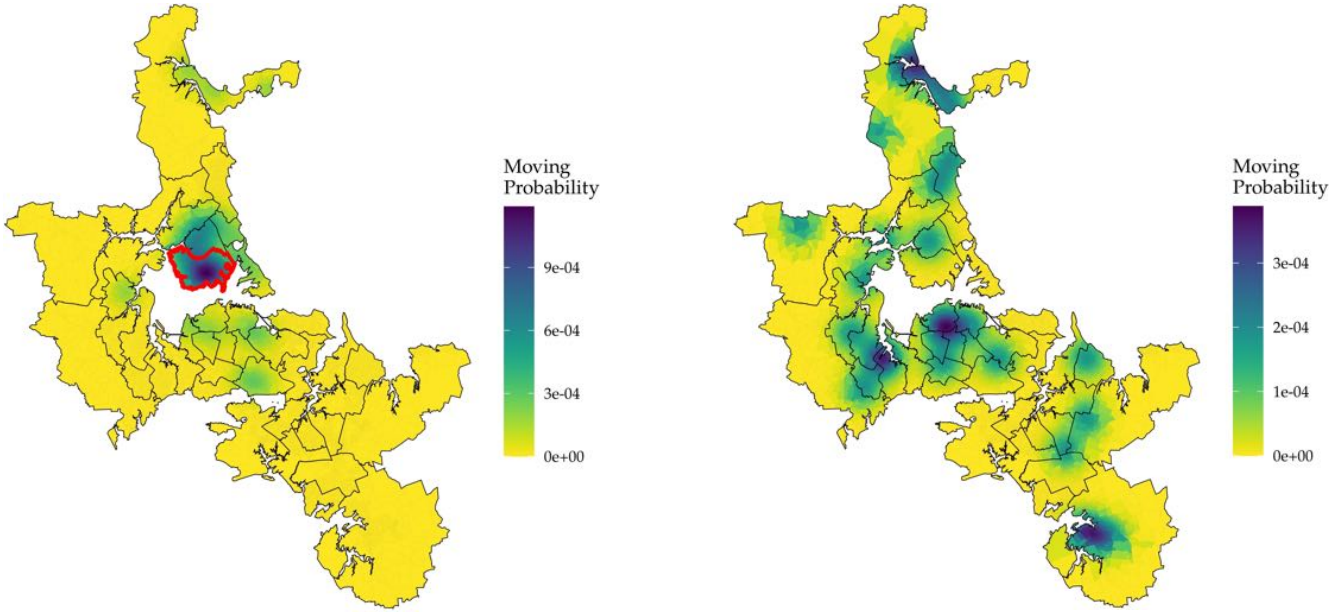

Note: The figures show the results of the spatial kernel density estimation (KDE) for two example state variables. In panel (a), the state has the current home location as the red highlighted area. In panel (b), the state has the current home location as the outside option. Probabilities are shown as the likelihood of moving to each Level 1 area unit (Figure A.16a).

Figure D.3: CCP versus Move Distance

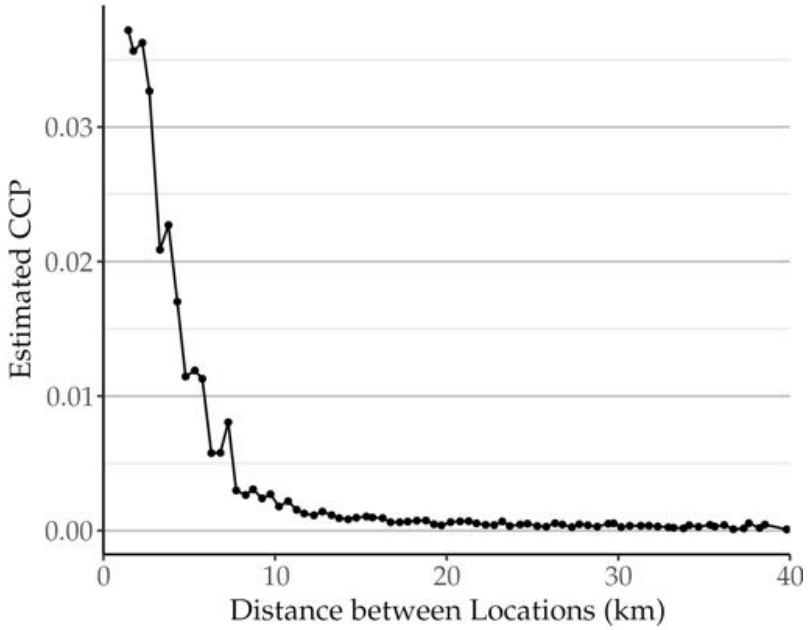

Note: The figure shows the average conditional choice probability (CCP) for a move, given the distance between the current home location and the new location. Averages are taken over state variable-action units (for actions that are moves inside the city). I group units using bins of 0.5 kilometer distances.

D.3 State Transition Estimation

Estimation Equation (14) requires knowing the state transition of the individual state variables (F_{dt}^k). Under the model, many of these variables transition deterministically and require no further estimation. However, there are two key stochastic processes: the timing of the second child's birth, and the school assignment process.

D.3.1 Timing of Second Child

Families enter the model already knowing the timing of their first child's birth (section 4.1). As I have restricted the estimation sample to people with at most two children, this means I only need to estimate the transition process for the second child's birth. Figure D.4 shows the empirical distribution (separately by ethnicity) of the difference in age between the first and second child. Across all ethnic groups, the modal age difference between the first and second child is two years. Asian and Polynesian families are more likely than white families to have an age gap greater than four years. Let the distribution shown in Figure D.4 be denoted as $gap(x, \psi)$, where $x \in \{0, 1, ..., 10\}$ is the size of the age gap. Then we

Figure D.4: Age Gap Distribution

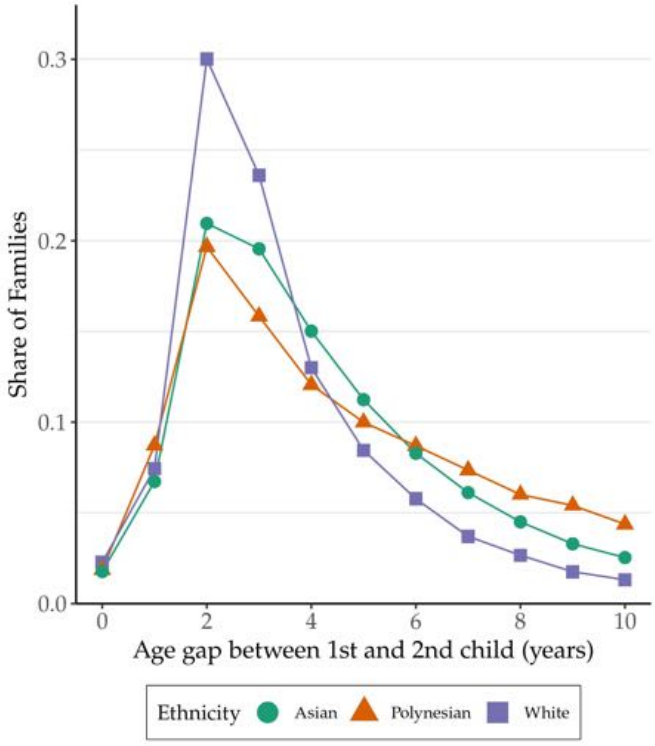

Note: The figure shows the distribution of the age gap between the first and second child (in years). I restrict to families with two children who have an age gap of no more than 10 years. The distribution is shown separately for each ethnic group.

can use it to fully express Equation (C.1) as:

$$Pr(T^b_{ic} = t) = B_\psi(C_i, C_{it}, a_{it,c-1}) = \begin{cases} \dfrac{gap\left(a_{it,c-1} - t, \psi\right)}{\left(\sum_{x=(a_{it,c-1}-t)}^{10} gap\left(x, \psi\right)\right)} & \text{if } C_{it} < C_i \\ 0 & \text{if } C_{it} = C_i \end{cases} \tag{D.1}$$

Intuitively, Equation (D.1) is calculating the probability of having a second child x years after the first child by dividing the probability mass shown in Figure D.4 over the cumulative probability of having an age gap of x to 10. This is effectively a normalization: when time is x periods after the first child's birth, it is no longer possible to have a second child with an age gap of less than x.

D.3.2 School Assignment

For the school distribution, I use the school records and match it with the family location to determine the probability of a child attending a school, conditional on their grade, home location and the calendar

year. To ensure that this captures the (new) school assignment distribution, $M_t^{New}(h_{it}, g_{itc})$, I estimate these probabilities only among children who are in the starting grade for each level (1, 7, and 9) or those who newly moved into the region.

Figure D.5 shows how the assignment probabilities vary on average with distance from home. These show all the expected patterns (for example, it is similar to the distribution shown in Figure A.10). In particular, the probability density is highest among schools that are closer to the family's home location. This distribution also implies that the expected distance to school is lowest for the primary school level and highest for high school level.[78]

Figure D.5: School Probability versus Distance

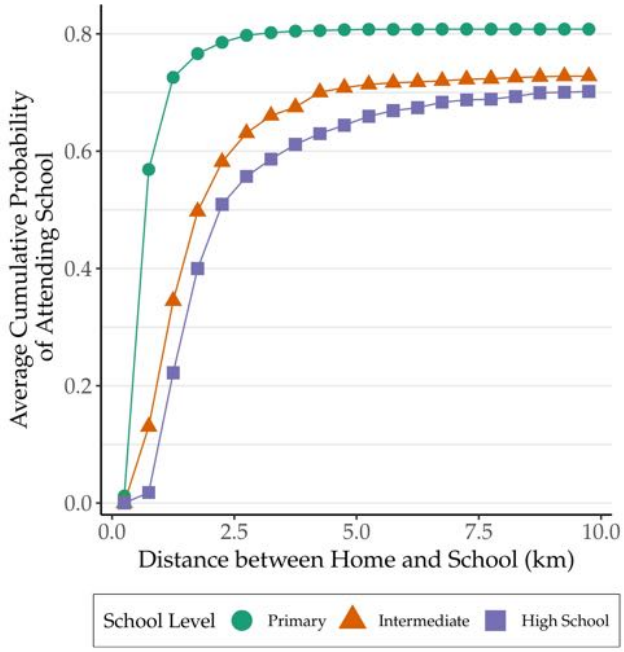

Note: The figure shows the average of cumulative distributions for each location-year. The vertical axis shows the probability of attending a school which has a distance from the home location less than or equal to the horizontal axis value. The probability is calculated separately by school level. The probabilities are then averaged over all location-years.

To calculate the reassignment probability ρ, I focus on two groups students: (1) students who stay in their current home, and (2) students who move to home within the same region as their previous home and do not start a new school level. According to Equation C.7, the probability of these children staying in the same school as last period (M_t^{Same}) is defined as:

$$M_t^{Same}(h_{it}, g_{itc}) \equiv (1-\rho) + \rho M_t^{New}(h_{it}, g_{itc}) = 1 - \rho\left(1 - M_t^{New}(h_{it}, g_{itc})\right)$$

[78] The probabilities do not reach 1 as the remainder is captured by the outside option school, for which there is no distance component as discussed in section 4.3.

$M_t^{Same}(h_{it}, g_{itc})$ can be estimated in the data as the proportion of these students who stay at the same school. After estimating $M_t^{New}(h_{it}, g_{itc})$ as described in Appendix C.1.4, I can then back out the reassignment probability as:

$$\rho = \frac{1 - M_t^{Same}(h_{it}, g_{itc})}{1 - M_t^{New}(h_{it}, g_{itc})}$$

I can do this to calculate a ρ value for each location, school, and year (I pool all grades within a school together). To reduce potential noise, I use the Level 4 locations (Figure A.16d), which are larger than the model locations (Level 3) still smaller than the regions (Level 5). The distribution of imputed reassignment is shown in Figure D.6. I take the median of the distribution as the estimated parameter value. This gives me $\rho = 3.69\%$.

Figure D.6: Reassignment Probability Estimates

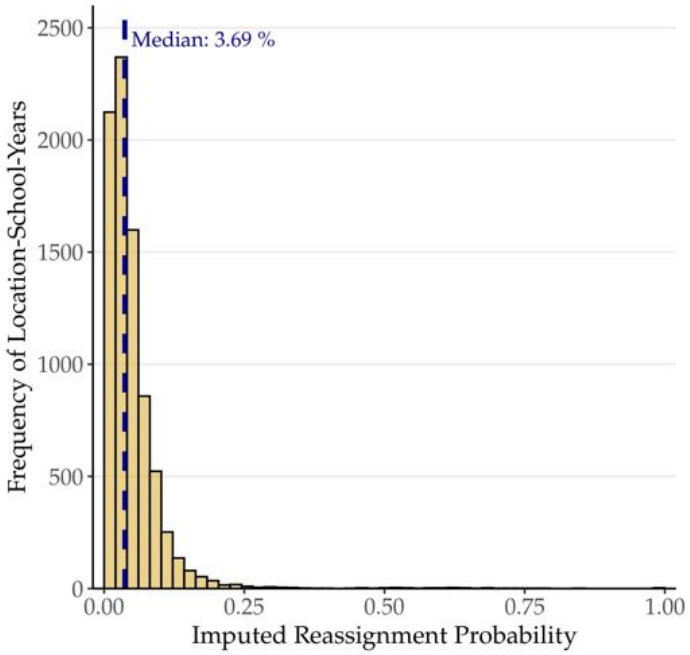

Note: The figure shows the distribution of estimated reassignment probabilities. Each unit is a location-school-year. The vertical line indicates the median value (3.69%) over all estimated values.

D.4 Building the Dataset

The regression equation (14) is at the state-action tuple level (k^*, t, d, d', d''). The state is an individual state k^* at a time t, where time captures the market states. The action tuple is a set of actions (d, d', d'') with two distinct time t actions (a baseline action $d_t = d$ and a counterfactual action $d_t = d'$) and a $t+1$ renewal action $d_{t+1} = d''$. Therefore, the data used to estimate this regression should be at the same level. In fact, the model implies that this equation holds for any state (k^*, t), any d_t pair (d, d'), and any (appropriately chosen) renewal action d''. However, an estimation that covers all possible state and

action combinations is not possible due to practical and computational limitations.

Instead, I construct the data in the following way. For each state, I consider two baseline actions d. These are selected as each state's two most likely actions (in terms of their CCP). Unsurprisingly, the $stay$ action always has the highest CCP among all states. The second action is chosen as the most likely out of $\mathcal{L} \cup \{0\}$. For the counterfactual action d', I choose all possible actions where $d' \neq d$.[79] Finally, for the renewal action d'', I choose the action that maximizes the joint probability over all possible renewal actions:

$$d'' = \underbrace{\underset{\tilde{d} \in \mathcal{D} \setminus \left(\{stay\} \cup \mathcal{R}(h_{it}) \cup \mathcal{R}\left(h'_{it}\right)\right)}{\operatorname{argmax}}}_{\substack{\text{Actions that are renewal actions} \\ \text{for both } d \text{ and } d'}} \left\{ \underbrace{F^k_{dt}(k^*)\,\mathbf{p}_{\tilde{d},t+1}}_{\substack{\text{Probability of action } d_{i,t+1}=\tilde{d} \\ \text{given initial action } d_{it}=d}} \times \underbrace{F^k_{d't}(k^*)\,\mathbf{p}_{\tilde{d},t+1}}_{\substack{\text{Probability of action } d_{i,t+1}=\tilde{d} \\ \text{given initial action } d_{it}=d'}} \right\}$$

In other words, d'' is chosen as the most likely renewal action given the state transitions after taking each action d and d'. This approach keeps the scope of the data to a manageable size, while still evaluating a broad range of possible actions. In practice, the actual dataset is smaller than this scope as I am not able to estimate all the CCPs for each state. However, the CCPs that are estimated are also, by definition, the most empirically relevant ones.

The final step is to determine the values of the regression equation components. The components of Equation (14) that require data analogues are determined as follows:

- $Y_{d,d',d'',t}$ is calculated according to Equation (12) using the estimated CCPs, state transitions, and the assumed discount factor β.
- SQ_{dt} is calculated as the weighted average of expected decile scores. The expectations are taken using the next period school admission transition probabilities, and the weighted according to the observed current period school attendance distribution. As I discuss in section 5.1, it is infeasible to include the attended school as a state variable. Therefore, we should interpret the CCPs as weighted probabilities given the underlying exogenous school distribution.
- DS_d is calculated using the average straight-line distance between the school location and the home location. The average is calculated as the mean distance between the school and each Level 1 centroid, weighted by the the proportion of the location living in each Level 1 unit. Similar to the school quality, this is then weighted by the school attendance distribution.
- $\ln r_d$ is calculated using the mean log sales and log rentals in the location-year. I make sales and rentals into comparable units by calculating the ratio of mean sale prices to mean annual rents in

[79] For the second baseline action $d \neq stay$, I also ensure that $d' \neq stay$ to avoid duplicating a symmetrical action pair. Given $|\mathcal{D}| = 44$, this process looks at $43 + 42 = 85$ possible (d, d') choices for each state.

each location-time and adjust the sales value using the corresponding ratio. See Figure D.7 for a distribution of the estimated ratios and further details. I find the median ratio to be 26.39.[80] I create a single price index by weighting the sale and rent values using their respective number of transactions in each location-year.

- TS_d^{φ} is calculated as the proportion of families living in each location-year that are of type φ.
- MC_d requires knowing the distance between two locations. I calculate this as the population-weighted pairwise distance of Level 1 units.[81]

Figure D.7: House Sales to Rents Estimated Ratios

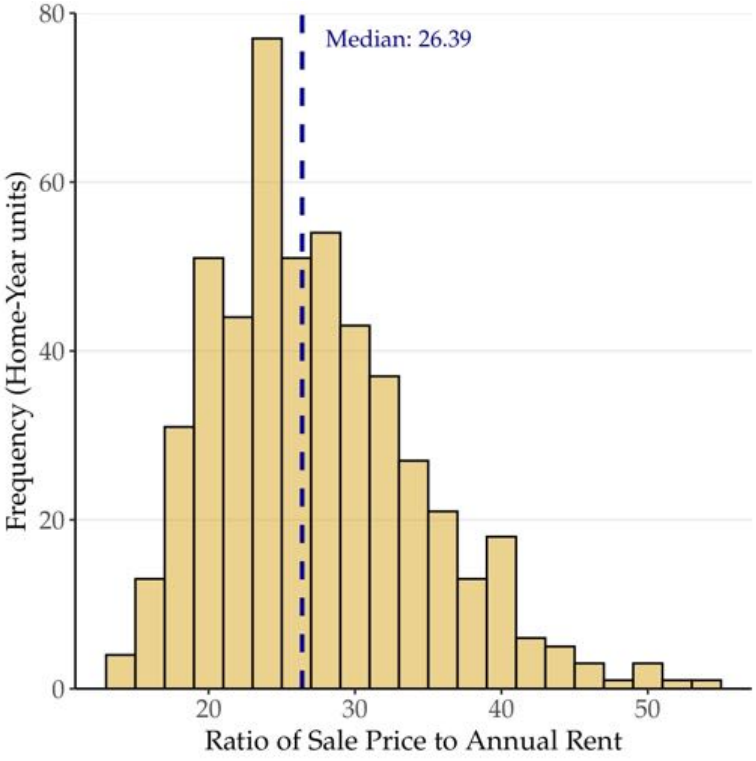

Note: The figure shows the distribution of the ratio between sales and annual rents in each location-year. Ratios are calculated for each Level 2 area (Figure A.16b) in each year. I use the mean sale price and mean annual rent (in real dollars) for the calculation. Then I calculate a location-year ratio (where a location is Level 3 unit, Figure A.16c) by weighting the Level 2 ratios using the number of transactions in the Level 2 area. The vertical line indicates the median value across all location-year ratios to be 26.39.

[80] Due to limited data, I only compare average prices without controlling for house characteristics. Bayer et al. (2007) do a similar adjustment using a hedonic regression to control for these factors. They find an average ratio of 264.1 between sales and monthly rents. This corresponds to a ratio of 22 in annual terms, which is comparable to my median figure.

[81] Concretely, for two locations h and j, the distance is calculated as:

$$\text{dist}(j,h) = \sum_{a \in Lv1(h)} \sum_{b \in Lv1(j)} \varpi_a \cdot \varpi_b \cdot \text{dist}_{Lv1}(a,b)$$

where a and b are the Level 1 units within h and j, respectively, and ϖ represents their population weights (which sum to 1 within each location). $\text{dist}_{Lv1}(a,b)$ is the straight-line distance between the centroids of a and b.

D.5 Identification Strategy Details

D.5.1 School Identification Strategy

Figure D.8: School Parameters Identification Strategy

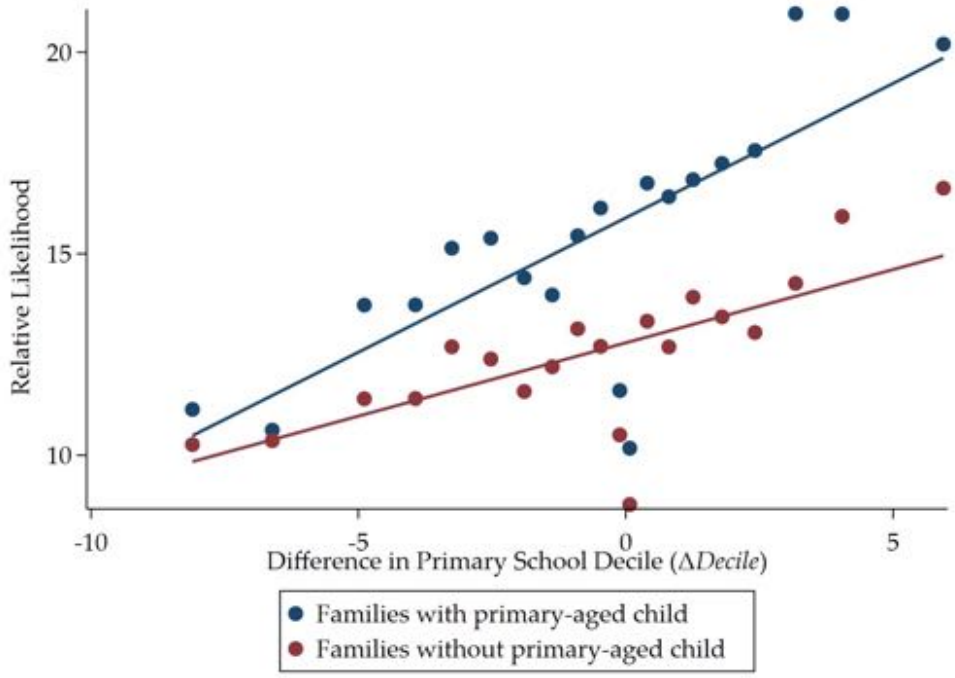

Note: The figure illustrates the motivation for the identification of the school parameters. I plot the relative likelihood of two action paths $(Y_{d,d',d'',t})$ against the difference in decile score experienced as a result of the two actions (d, d'). I average the relative likelihood into binned values of the decile score difference. This relationship is plotted for families with a child in primary school (blue) and for families with a child was not yet started primary school (red).

Figure D.8 gives a graphical illustration of how the school identification strategy works. I compare the relative likelihood value ($Y_{d,d',d'',t}$) against the expected difference in primary school decile score resulting after the two actions (d and d'). I plot this for families with children in primary school and for families with children who have not yet started school. The positive correlation shows that families with school-aged children are more likely to move to areas with relatively higher decile scores. However, we cannot attribute all of this to school quality alone as these areas may also have other desirable characteristics. This is suggested in the figure as families without school-aged children also have a higher likelihood of moving to a place with higher decile quality, even though they receive no benefit from schools in the given time period.[82] Therefore, by differencing the relative likelihoods across the two groups, I can isolate the effects of schooling from other neighborhood effects. In other words, we can use the families without school-aged children to control for the unobserved neighborhood characteristics. The difference

[82]This is not saying that families do not *ever* benefit from schools. The ECCP methodology means I am only evaluating the *flow* utility received from each location.

between the two lines in Figure D.8 captures the left hand side of Equation (15). Note that this figure is showing the relationship only for (primary) school quality, while the estimation of Equation (15) is simultaneously capturing the effects for the other school variables (distance, moving costs, and outside options), and for all three school levels.

To make the comparison as best as possible, I compare families who both have one child born. For families with a primary school child, I compare to families with a child who has not yet started primary school. For families with a high school child, I compare to families with a child who has already completed high school. For families with an intermediate school child, I take an average of the previous two control groups (not yet started primary school and completed high school).

Evaluating Equation (14) for families without school-aged children gives the following likelihood value:

$$\begin{aligned} Y_{d,d',d'',t}\{\text{Non-school-aged}\} =& \alpha_\psi^r \left(\ln r_{dt} - \ln r_{d't}\right) + \sum_{\varphi \in \Psi} \alpha_\psi^\varphi \left(TS_{dt}^\varphi - TS_{d't}^\varphi\right) \\ &+ \left(\delta_{\psi d} - \delta_{\psi d'}\right) + \alpha_\psi^\tau \left(\tau_d - \tau_{d'}\right) \\ &+ \left(MC_d - MC_{d'}\right) \\ &+ \beta \left(MC_{d,d''} - MC_{d',d''}\right) + \epsilon_{dd't} \end{aligned} \tag{D.2}$$

D.5.2 Time Differencing

After estimating the parameters of Equation (15), I can create the predicted values using the estimated coefficients. Let $\widetilde{Y}_{d,d',d'',t}$ represent the residuals after removing the school-related components from $Y_{d,d',d'',t}$:

$$\begin{aligned} \widetilde{Y}_{d,d',d'',t} = Y_{d,d',d'',t} &- \left(\widehat{SQ}_{dt} - \widehat{SQ}_{d't}\right) - \left(\widehat{DS}_d - \widehat{DS}_{d'}\right) - \left(\widehat{SMC}_{dt} - \widehat{SMC}_{d't}\right) \\ &- \left(\widehat{OOS}_{dt} - \widehat{OOS}_{d't}\right) + \beta \left(\widehat{SMC}_{d,d'',t} - \widehat{SMC}_{d',d'',t}\right) \end{aligned}$$

where the hats $\widehat{\cdot}$ represents the values calculated using the estimated coefficients. This simplifies to:

$$\begin{aligned} \widetilde{Y}_{d,d',d'',t} =& \alpha_\psi^r \left(\ln r_{dt} - \ln r_{d't}\right) + \sum_{\varphi \in \Psi} \alpha_\psi^\varphi \left(TS_{dt}^\varphi - TS_{d't}^\varphi\right) \\ &+ \left(\delta_{\psi d} - \delta_{\psi d'}\right) + \alpha_\psi^\tau \left(\tau_d - \tau_{d'}\right) \\ &+ \left(MC(d) - MC(d')\right) + \beta \left(MC_d\left(d''\right) - MC_{d'}\left(d''\right)\right) + \epsilon_{dd't} \end{aligned} \tag{D.3}$$

Now I consider a state-action tuple (k^*, t, d, d', d'') and another state-action tuple (k^*, s, d, d', d''), where $t > s$. These units are similar on all dimensions except the calendar year. Their corresponding residuals

can be expressed as consider $\widetilde{Y}_{d,d',d'',t}$ and $\widetilde{Y}_{d,d',d'',s}$, respectively. Differencing this over time gives:

$$\begin{aligned}
\Delta_{ts}\widetilde{Y}_{d,d',d''} =& \widetilde{Y}_{d,d',d'',t} - \widetilde{Y}_{d,d',d'',s} \\
=& \alpha^r_\psi \left((\ln r_{dt} - \ln r_{d't}) - (\ln r_{ds} - \ln r_{d's}) \right) \\
&+ \sum_{\varphi \in \Psi} \alpha^\varphi_\psi \left(\left(TS^\varphi_{dt} - TS^\varphi_{d't} \right) - \left(TS^\varphi_{ds} - TS^\varphi_{d's} \right) \right) \\
&+ \left(\left(\delta_{\psi d} - \delta_{\psi d'} \right) - \left(\delta_{\psi d} - \delta_{\psi d'} \right) \right) + \alpha^\tau_\psi \left((\tau_d - \tau_{d'}) - (\tau_d - \tau_{d'}) \right) \\
&+ \left(\left(MC(d) - MC(d') \right) - \left(MC(d) - MC(d') \right) \right) \\
&+ \beta \left(\left(MC_d \left(d'' \right) - MC_{d'} \left(d'' \right) \right) - \left(MC_d \left(d'' \right) - MC_{d'} \left(d'' \right) \right) \right) \\
&+ \epsilon_{dd't} - \epsilon_{dd's} \\
=& \alpha^r_\psi \left(\Delta_{ts} \ln r_d - \Delta_{ts} \ln r_{d'} \right) + \sum_{\varphi \in \Psi} \alpha^\varphi_\psi \left(\Delta_{ts} TS^\varphi_d - \Delta_{ts} TS^\varphi_{d'} \right) + \Delta_{ts} \epsilon_{dd'}
\end{aligned}$$

The interpretation of Equation (14) was to compare the difference in the utility of location A versus location B. Now that I am differencing over time, I am calculating how that difference in utility has *changed* over time and seeing whether that can be explained by how much the neighborhood variables have themselves *changed* over time. There may still be omitted variable bias as these prices could be correlated with other neighborhood amenities that have also changed over time. For this reason, I need instruments to address this endogeneity.[83] I expect the primary source of endogeneity to be the unobserved neighborhood characteristics (ξ). However, one could also be concerned that the expectational errors could also bias the estimates. As with the other terms, the expectational error enters as the difference of two actions (see Equation (C.21) in Appendix C.4). Therefore, the endogeneity concern would be if expectational errors vary significantly after living in a different location for one year – and that this effect on expectations is changing over time. This seems unlikely and a second-order concern as compared to the unobserved neighborhood characteristics.

Using an IV regression in this case requires finding three instruments as there are three endogenous variables. One of these variables is the difference in house price ($\Delta_{ts} \ln r_d - \Delta_{ts} \ln r_{d'}$). The other two are differences in type shares for Polynesians and Asians, ($\Delta_{ts} TS^\varphi_d - \Delta_{ts} TS^\varphi_{d'}$), as the ethnic share for whites is excluded to avoid multicollinearity. When I apply the instruments, I transform them in the same way as the endogenous variables. This means that they are differenced across locations (capturing the difference between d and d') and then differenced across time (capturing the difference between t and s).

[83] Note that a similar strategy to identifying the school variable coefficients would not work, as families (within a type) value prices and type shares in the same way.

D.5.3 Price Instrument Construction

The chosen instrument for price is the average price in non-neighboring locations. In this context, since the regression equation is a difference between two locations, I need to ensure that it is non-neighboring to *both* locations. If not, then one location's instrument could be correlated with the other location's price variable, which would violate the exclusion restriction. The intuition for constructing this instrument is shown using two example locations in Figure D.9. There are two locations in the unit of analysis (locations 1 and 2). I exclude all areas within a 10 kilometer buffer of each of these areas.

To generate variation among the remaining set of locations, I assign each location a weight according to its inverse distance to the location of interest. This means that prices in areas closer to location 1 should have a greater effect on location 1 prices than they do on location 2 prices. Generating this variation is critical. Let $\bar{r}_{-N(j,j')}$ be the average non-neighboring price for a location pair j and j'. If this were a simple average over all locations, then $\bar{r}_{-N(j,j')} = \bar{r}_{-N(j',j)}$. This means differencing equations across the two locations will eliminate the price instrument component, which means its coefficient could not be identified. By using a weighted average (with different weights for each location), this avoids the issue.

I create each location's distance-weighted non-neighboring price index and use that as the instrument. With this approach, I construct a different value of the instrument for each location pair in each year.

D.5.4 Ethnicity Instrument Construction

In 1987, New Zealand changed its immigration policies to eliminate preferences for certain nationalities, resulting in increased immigration during the 1990s and 2000s from non-traditional countries, especially east Asia (Grbic et al., 2010). As the Asian and Polynesian populations in Auckland are highly tied to immigration patterns, these new immigrants were likely to locate themselves in areas with higher shares of their own group, which motivates the use of the shift-share instrument.

To find the initial shares, I use aggregate statistics from the 1986 census, which is the earliest data that I could find at the sub-city level (and, importantly, this is prior to the 1987 immigration reform). The city had a remarkably different ethnic make-up then: in 1986, 76.2% of Auckland's population was white and only 2.8% was Asian.

Using the census statistics, I calculate the population of each ethnic group living in each of the city's locations as of 1986 ($pop^{\psi}_{j,1986}$). I divide this by the national population of ψ in 1986 ($natpop^{\psi}_{1986}$) to calculate the share of the national population of group ψ that is living in location j as of 1986. This represents the baseline shares in the past settlement instrument. For each time period t, I scale the baseline share by the national population of ψ as of time t ($natpop^{\psi}_{t}$). This step (Equation D.4) generates

Figure D.9: Price Instrument Example

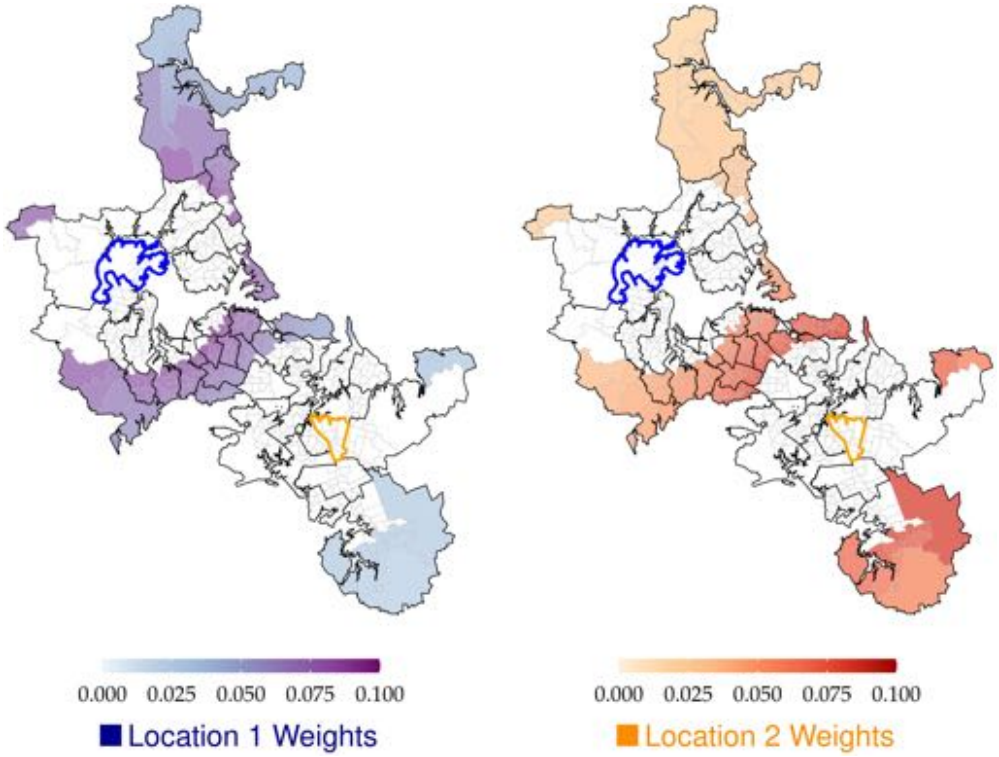

Note: The figure illustrates the price instrument calculation. Given a location pair (i.e. the locations after $d_t = d$ or d'), I find all Level 2 areas (Figure A.16b) that are more than 10 kilometers away from both locations. Among those areas, their prices are weighted using their average inverse distance to the locations of interest. The left map plots the value of the weights for location 1 (highlighted with a blue border). The right map plots the value of the weights for location 2 (highlighted with an orange border). Excluded areas with no weight are unshaded.

a simulated population of each ethnic group ψ in location j at time t.

$$simpop_{jt}^{\psi} = natpop_{t}^{\psi} \times \underbrace{\frac{pop_{j,1986}^{\psi}}{natpop_{1986}^{\psi}}}_{\substack{\text{Share of New Zealand's } \psi \text{ population} \\ \text{living in location } j \text{ in 1986}}} \tag{D.4}$$

The simulated population ($simpop_{jt}^{\psi}$) represents the expected ethnic population in each location if the annual national population continued to be distributed according to the 1986 population distribution. From this, I can then calculate the implied simulated ethnic share as:

$$simshare_{jt}^{\psi} = \frac{simpop_{jt}^{\psi}}{\sum_{\varphi \in \Psi} simpop_{jt}^{\varphi}} \tag{D.5}$$

By using the 1986 ethnic shares and shifts in the national population, this instrument should capture changes in the ethnic shares that are independent of the unobservable neighborhood utility ξ. This is for two reasons. First, the national population growth should be independent of local time-varying shocks

occurring in Auckland's neighborhoods. Second, even if this instrument were correlated with unobserved neighborhood characteristics, then it would be primarily correlated with neighborhood utility as of 1986 (but not with neighborhood utility from the sample period of 2008-2019).

A potential threat to the identification strategy is if ξ exhibits serial correlation such that unobserved neighborhood utilities in 2008-2019 are still correlated with their 1986 values. There are historical and contextual reasons why this is unlikely. First, this is a time span of 22 years, during which there was rapid development of the city (Friesen, 2009; Auckland Regional Council, 2010).[84] Second, the change in immigration policy resulted in a observably different set of immigrants arriving into New Zealand as compared to past migrants (Ho, 2015). Therefore, even if the unobserved characteristics remained the same, the ethnic group's aggregate preferences towards neighborhood amenities is unlikely to have matched the preferences of those already living in the city as of 1986. Figure D.10 shows that while the simulated share is correlated with the actual shares, this correlation appears to be getting weaker over time. This supports the idea that neighborhood effects in the past are likely to be less of an endogeneity concern over time.

Figure D.10: Ethnicity Instrument

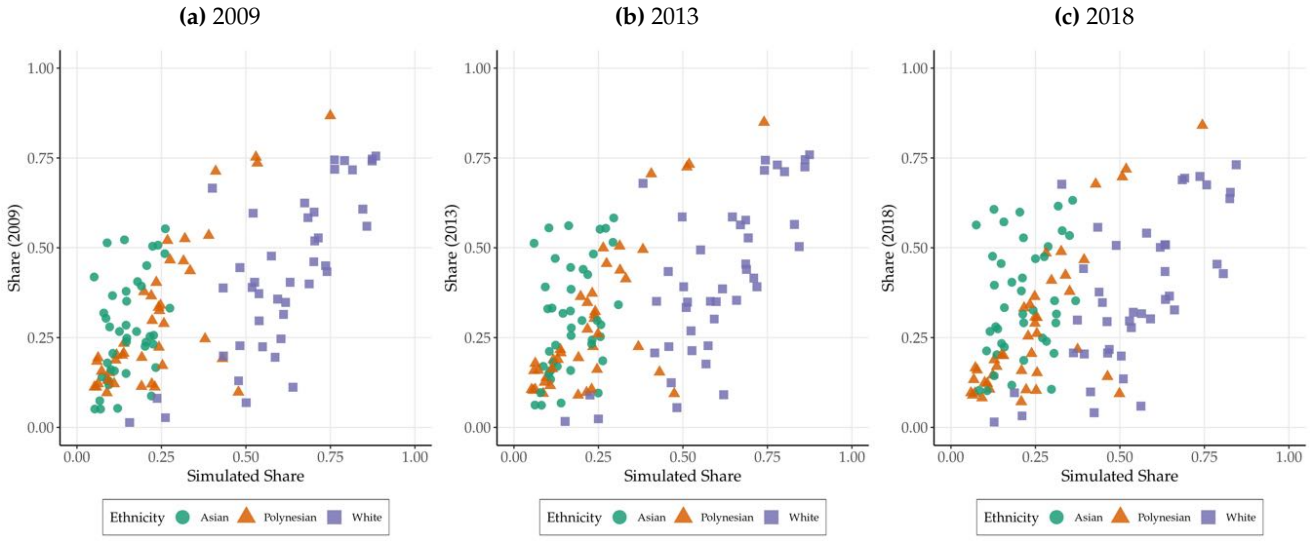

Note: The figure compares the simulated ethnic share and the observed ethnic share for each location-ethnic group. The panels represent the calculations of the shares (both simulated and observed) at different years.

Note that, as the regression employs a differencing strategy, I also need to difference the ethnicity instrument (Equation D.5) to apply it in the estimation. Figure D.11 shows how this works. In each graph, I plot for each location and each time the actual ethnic share versus the simulated ethnic share. As already shown, the simulated and actual shares are highly positively correlated. However, in this context, I am differencing over time, which may generate a different relationship. In the case for Asian share

[84] To give a sense of the rapid development, 76% of the population growth in New Zealand between 1991 and 2001 occurred in Auckland. This population growth affected existing residential areas rather than solely expanding the urban region. In particular, over this same period (1991-2001), 52% to 62% of annual residential growth in Auckland took place in the existing urbanized areas (Auckland Regional Council, 2010).

(Figure D.11a), we can see that – within a location – a positive increase in the simulated share over time is correlated with a positive increase in actual share. However, in the case of Polynesian share (Figure D.11b), the relationship is the opposite. As the simulated share increases (within the same location over time), the observed ethnic share generally decreases.

Figure D.11: Ethnicity Instrument Over Time

(a) Share Asian

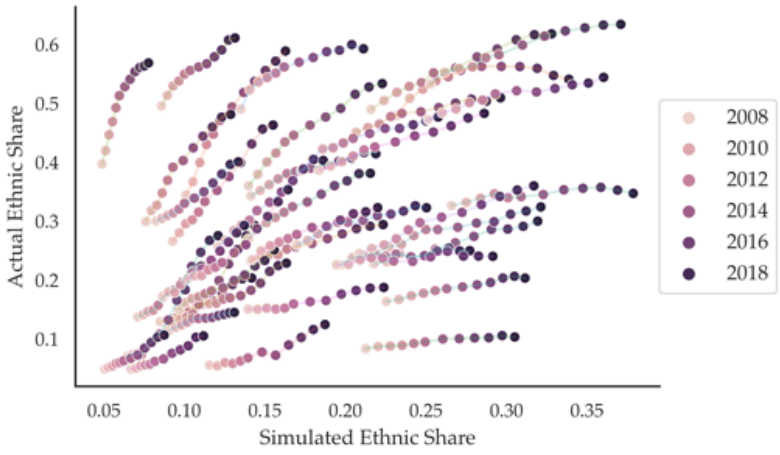

(b) Share Polynesian

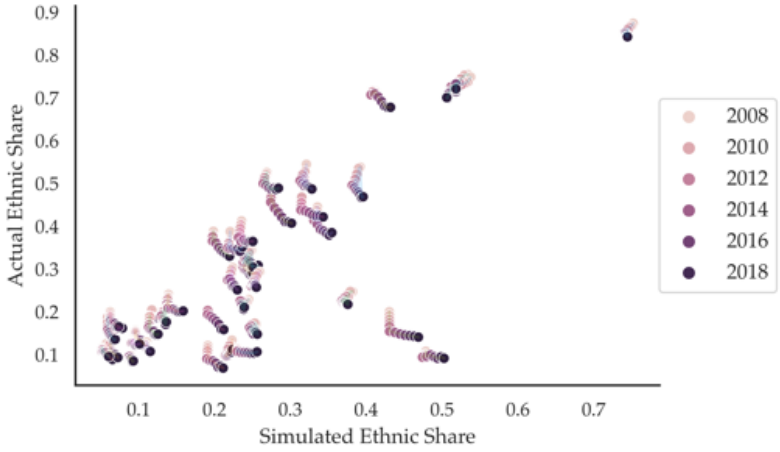

Notes: The figures show the scatter plot of actual ethnic share versus the simulated ethnic share for each location-year. Markers for the same location are connected by a line. The hue of each marker represents the year, ranging from 2008 to 2019 (later years are shaded darker).

D.6 Regression Tables

- Table D.2 shows the estimation of the school parameters (section 5.4)
- Table D.3 shows the IV estimation of the neighborhood parameters (section 5.5)
- Table D.4 shows the estimation of the moving and location parameters (section 5.6)

Table D.2: School Parameter Estimation

	School Level	(1) White	(2) Polynesian	(3) Asian
Decile Score	High School	0.273***	0.262***	0.176*
		(0.0652)	(0.0761)	(0.101)
	Intermediate	−0.0255	0.552***	0.382***
		(0.0980)	(0.103)	(0.108)
	Primary	0.126***	0.394***	0.304***
		(0.0473)	(0.0464)	(0.0764)
School Moving Cost	High School	−2.392***	−1.751***	−1.990***
		(0.233)	(0.265)	(0.323)
	Intermediate	−2.615***	−1.946***	−2.306***
		(0.439)	(0.409)	(0.449)
	Primary	−3.822***	−2.661***	−3.117***
		(0.293)	(0.260)	(0.333)
Distance (km)	High School	−0.677***	0.266 ***	0.374
		(0.226)	(0.469)	(0.424)
	Intermediate	−1.222*	−0.420	−2.302***
		(0.642)	(0.506)	(0.554)
	Primary	−0.264	1.799***	−0.440
		(0.206)	(0.328)	(0.449)
Distance Squared (km^2)	High School	0.0557**	−0.101	−0.00986
		(0.0294)	(0.102)	(0.0855)
	Intermediate	−1.222*	−0.420	−2.302***
		(0.642)	(0.506)	(0.554)
	Primary	0.0516***	−0.273***	0.178*
		(0.0191)	(0.0611)	(0.104)
Outside Option	High School	−1.076	0.979	1.821
		(0.908)	(0.816)	(1.145)
	Intermediate	0.392**	0.505***	0.525***
		(0.166)	(0.148)	(0.112)
	Primary	−1.038	−4.135**	−1.060
		(1.614)	(1.645)	(1.774)
N		1,713,570	2,027,559	2,112,000
R^2		0.027	0.030	0.027

Note: The table shows the estimates for the school utility parameters, each of which can differ by school level (Equation (15)). Each column represents the estimation results for the ethnic group listed at the top. The unit of observation is a state s_{it}-action tuple (d, d', d'') pairing. The outcome variable is the difference in relative log likelihood of two action paths $(\Delta Y_{d,d',d'',t})$, where the difference is taken between a state representing a family with one child in school and states representing otherwise similar families with one child not in school. The values of $Y_{d,d',d'',t}$ corresponding to the non-school-aged child states are averaged into one value. Standard errors, shown in parentheses, are clustered at the home location level. In accordance with Stats NZ confidentiality rules, the number of observations have been randomly rounded to a multiple of 3 (RR3). Significance stars represent the following p-values: $^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$.

Table D.3: Neighborhood Variables Parameter Estimation

(a) First Stage Estimates

	(1) Price	(2) Polynesian Share	(3) Asian Share
Non-Neighboring Price	0.586***	−17.86***	56.54***
	(0.0530)	(2.099)	(6.340)
Simulated Polynesian Share	−0.00874***	−0.636***	1.954***
	(0.00155)	(0.0482)	(0.147)
Simulated Asian Share	0.00602***	−0.119***	0.204***
	(0.000670)	(0.0197)	(0.0663)
Constant	0.0000336	0.0389**	−0.165***
	(0.000534)	(0.0191)	(0.0540)
N	77,808	77,808	77,808
R^2	0.066	0.102	0.147
F	118.8	74.93	96.84

(b) Second Stage Estimates

	(1) White	(2) Polynesian	(3) Asian
Log Annual Price	−24.86***	−37.69**	−28.54***
	(5.647)	(17.89)	(4.662)
Polynesian Share (%)	−1.351	−1.927*	−0.825*
	(1.333)	(1.095)	(0.462)
Asian Share (%)	−0.265	−0.0890	−0.278***
	(0.422)	(0.159)	(0.0976)
Constant	−0.354***	−0.113*	−0.444***
	(0.0549)	(0.0657)	(0.0534)
N	77,808	72,228	69,114

Note: The table shows the results of estimating Equation (16) by instrumental variables. Prices are the log of the annual price index. Share variables and their instruments are measured in percentage points (0-100). The unit of observation is a neighborhood tuple $((h_t, j_t), (h_s, j_s))$ for two locations (h and j) being compared at two points in time (t and s). In the first stage, each column represents the outcome variable of the regression (the endogenous variable in the second stage regression). Results shown only for white families. Other ethnicity's first stage results are similar and available upon request. In the second stage, the outcome variable is the time-difference in residualized relative log likelihood of two action paths $(\Delta_{ts}\widetilde{Y}_{d,d',d''})$ averaged over states with the location outcomes (h, j) in period $t+1$ and $s+1$. Each column represents a separate regression by ethnic group. In accordance with Stats NZ confidentiality rules, the number of observations have been randomly rounded to a multiple of 3 (RR3). Significance stars represent the following p-values: $^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$.

Table D.4: Moving and Location Variables Parameter Estimation

	(1) White	(2) Polynesian	(3) Asian
Tenure	0.0756**	0.0949**	−0.0126
	(0.0343)	(0.0389)	(0.0292)
Max Tenure Indicator	0.562***	0.297**	−0.176*
	(0.147)	(0.111)	(0.0913)
Moving Cost: Distance	−0.765***	−0.662***	−0.631***
	(0.00959)	(0.00800)	(0.0112)
Moving Cost: Fixed	−2.811***	−1.479***	−3.767***
	(0.121)	(0.129)	(0.177)
Moving Cost: Outside	−2.546***	−2.470***	−3.504***
	(0.162)	(0.272)	(0.158)
Location Fixed Effects	Yes	Yes	Yes
N	19,003,770	16,216,170	17,524,269
R^2	0.983	0.994	0.961

Note: The table shows the parameter estimates for the moving and location variables. The outcome variable is a residualized relative log likelihood of two action paths. The unit of observation is a state $\mathbf{s}_{it}$-action tuple (d, d', d'') pairing. The maximum tenure value is 5 years. The distance moving cost is per kilometer, while the other two are fixed costs. Location fixed effects are included in each regression. Each column represents a separate regression by ethnic group. In accordance with Stats NZ confidentiality rules, the number of observations have been randomly rounded to a multiple of 3 (RR3). Significance stars represent the following p-values: $^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$.

E Discussion Appendix

E.1 School Quality: Reduced Form Normalization

To make the reduced form estimates into a per-child-per-year valuation (and therefore comparable to the structural estimates), we would need to account for the family's children. However, I am unable to identify the families who generate the house sales/rentals in the BDD analysis.

As an alternative, I use the family panel data and identify households who move into the BDD sample areas (i.e. the areas close to the boundaries) during the sample years. I cannot distinguish whether these people are renters or home-owners, but this approach does allow me to generate aggregate statistics about their children. I calculate each family's present discounted value of school years for each school level. For example, consider a family i with C_i children who move into a primary school sample area at time t_i^*. I calculate their present discounted value (PDV) of primary school years as:

$$PDV_i^{Prim} = \sum_{t=t_i^*}^{\infty} \sum_{c \in C_i} \beta^{t-t_i^*} \mathbb{1}\left\{g_{itc} \in [1,6]\right\} \tag{E.1}$$

The interpretation of Equation (E.1) is that this is the number of years their children will spend in primary school, summed over all their children and discounted by the discount rate β. I can do an analogous PDV calculation for intermediate $\left(PDV_i^{Int}\right)$ and high school $\left(PDV_i^{HS}\right)$. Among the new movers subpopulation, I find that the mean PDVs for each school level are: $PDV^{Prim} = 3.53, PDV^{Int} = 1.09, PDV^{HS} = 2.69$.[85] These PDVs reflect the "amount" of school quality (within a level) that the family expects to experience for the remainder of time. The PDVs being greater than 1 show that the reduced form estimates are not comparable to the structural model.

I use these figures to rescale the reduced form estimates to generate a per-child-per-year valuation.[86] However, this rescaling comes with an assumption that families do not move again after they moved into the BDD sample area. With this assumption, the PDV therefore captures the total school years the family expects to enjoy in their new home. I cannot calculate the average number of years families expect to stay in the house as the future cannot be observed in the data.[87] This is a strong assumption – indeed I find that 26% of the new mover households move again after only one year (which is consistent with the broader population, as shown in Figure C.7). To make it more plausible, I use the reduced

[85] These figures represent the mean PDV of a school level v for families who move into areas that are in the analysis sample for school level v. Note that the mean number of children already born when these families move into the BDD areas is 0.7, while the mean number of children they will have is 1.8. The average age (at the time of the move) among children who are born is 8.8. The PDV captures the school needs arising from the currently born children as well as the children that will be born in the future.

[86] I divide the reduced form estimate for school level v by the mean PDV value PDV^v.

[87] For example, consider a family who moves into the sample area in 2015 but has not left by 2019 (when my data ends). It would be incorrect to say that their housing tenure is 4 years, since that is driven by data censoring. Given this information, I cannot conclude what their full housing tenure will be.

form estimates from the sales data (column 2 of Table 1a). This is because it is more likely that those who purchase houses will stay in those locations for a long number of years. This also makes my benchmarking more relevant to the broader literature, as the vast majority of studies use house sales data (rather than rentals) when conducting their BDD analysis.

E.2 School Moving Cost: Example

Consider a white family with a first grader who still has 5 years of primary school to complete. Their WTP for an additional primary school decile score was estimated to be 0.508% per child per year. Their WTP for the school moving cost at primary level was estimated at –14.25%. Using the discount rate of $\beta = 0.9$, a back-of-the-envelope calculation shows that required increase in decile score (for the remaining primary school years) would have to be:

$$\Delta\text{Decile} = \frac{-\text{School Moving Cost}}{\text{Discounted Future Value of Additional Decile}} = \frac{14.25}{\sum_{t=0}^{4} \beta^t \times 0.508} \approx 6.85$$

The family would need to change to a school with a 6.85 higher decile score to compensate for their school moving cost. This change in decile score can be seen as an increase of more than 2 standard deviations in the underlying raw measure. Moreover, this calculation is keeping all else equal - in reality moving from (say) a decile 3 to a decile 10 school would mean moving to a substantially more expensive neighborhood. Therefore, it appears unlikely that parents can simply offset these school moving costs with higher school quality.

E.3 Counterfactuals: Simplifications

A limitation of the ECCP methodology is that is requires researchers to impose strong assumptions when conducting counterfactuals (Kalouptsidi et al., 2020). This is because the approach does not require a full model to be specified. Moreover, the computational burden associated with solving dynamic models – much of which were circumvented by using the ECCP estimator – are still present for the model simulations. For these reasons, I simplify the simulations by transforming the problem into a static model. As families in my simulations are myopic, this avoids having to forward simulate to calculate value functions.

To account for the fact that my estimates are from a dynamic model while the simulation is a static model, I rescale the location fixed effects to match empirical population patterns. For each ethnicity ψ and family size (number of born children C), I find a vector of time-varying location scales $(\zeta_{1t}^{\psi C}, ..., \zeta_{Lt}^{\psi C})$ to multiply to each location's estimated fixed effects. For a given period t, I solve for these scales using simulated method of moments to match the empirical residential distribution of each subgroup (ethnicity and family size cells). Starting at a period t, for a given choice of scales, I simulate the next period

$t + 1$ outcome using the neighborhood characteristics observed in the data and estimated parameters. I then calculate the proportion of the subgroup living in each location per the model simulation and compare this to the empirical analogs. I solve for the scales using a numerical optimization solver, where the objective function is to minimize the difference between the simulated and empirical shares. I repeat this for three periods (2008, 2012, and 2019) and extrapolate using a cubic spline interpolation for the remaining periods. This rescaling accounts for dynamic forces not captured by a static model, such as changing family needs as children are born and unobservable time-varying neighborhood characteristics. It should also account for the fact that the moving costs estimates are experienced in only one period while the benefits of a move are enjoyed for multiple periods.

After solving for these scales, I simulate the market where each period I solve for prices such that the total demand of housing units in each location is equal to the available supply of housing units in each location. I treat supply as exogenous and fixed in this model. I set supply for each location-year equal to the total number of families observed to be living in that location during that time.

I impose an additional simplification where I only allow a fraction of families to move each period. These families are randomly chosen at the start of each period (and are simply given the option of changing their location, though they can choose to stay). The remainder of families must stay in their current home. I calibrate this fraction to be 25% of families. This approach is a reasonable approximation of capturing that families do not often change their homes (e.g. it is a similar approach to modeling sticky prices in Calvo models of price rigidity). I impose this simplification because I find that without it, families move much more than expected (i.e. relative to the data). This issue is amplified each period as moving generates a positive feedback loop. As families move, the ethnic composition and prices in locations change, which results in more families moving the following period.

E.4 Supporting Figures

Figure E.1: Willingness to Pay for Other Utility Components

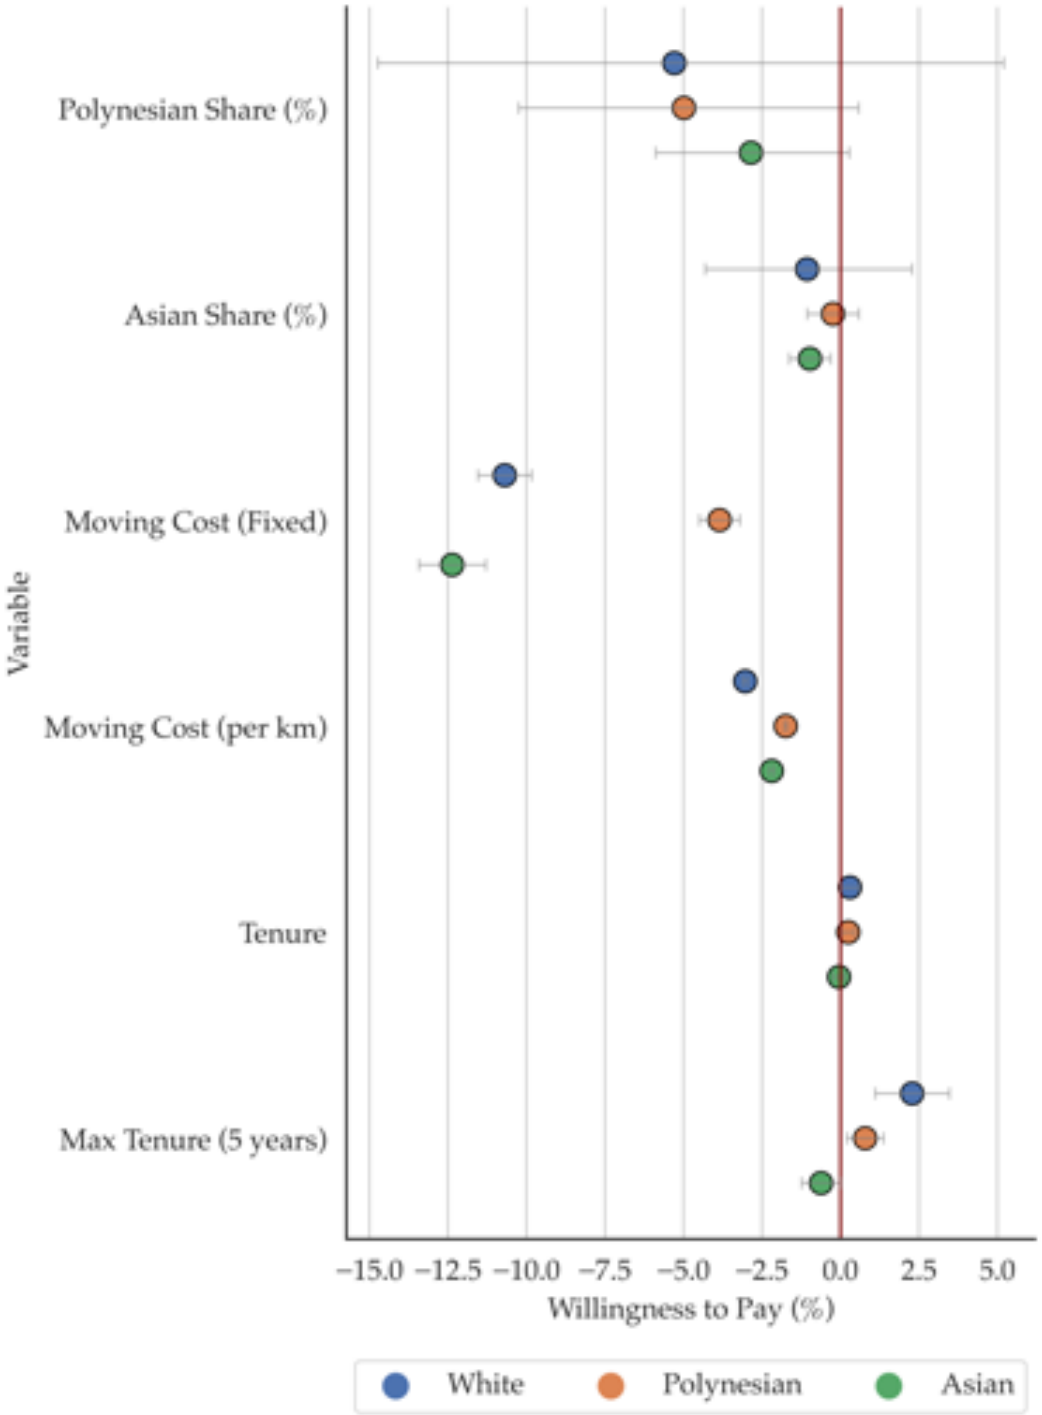

Note: The figure shows the willingness to pay estimates for components of location utility. The results use the structural model estimates (Tables D.3b and D.4) and scale by the ethnic group's price sensitivity estimate (Table D.3b). 95% confidence intervals are plotted for each estimate.

Figure E.2: Change in Home Location's Polynesian Share Among Polynesian Movers

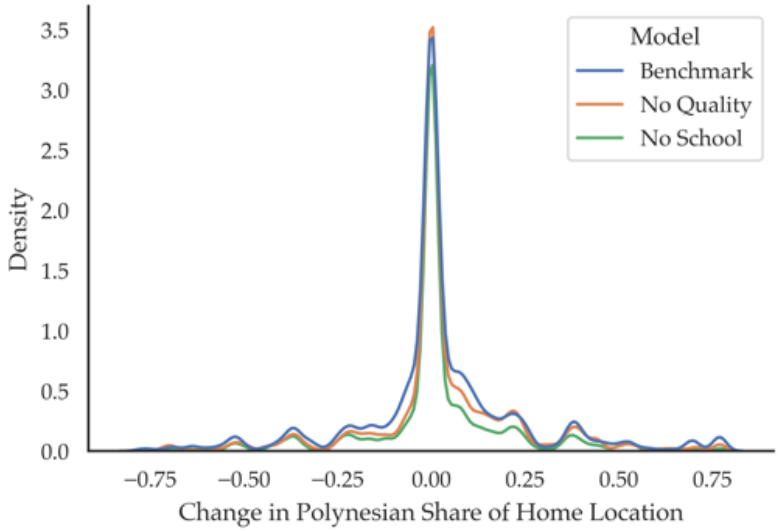

Note: See notes under Figure 9. This figure shows the equivalent statistic for Polynesian families who move and the change in the Polynesian share of their home locations.

Figure E.3: Change in Home Location's Asian Share Among Asian Movers

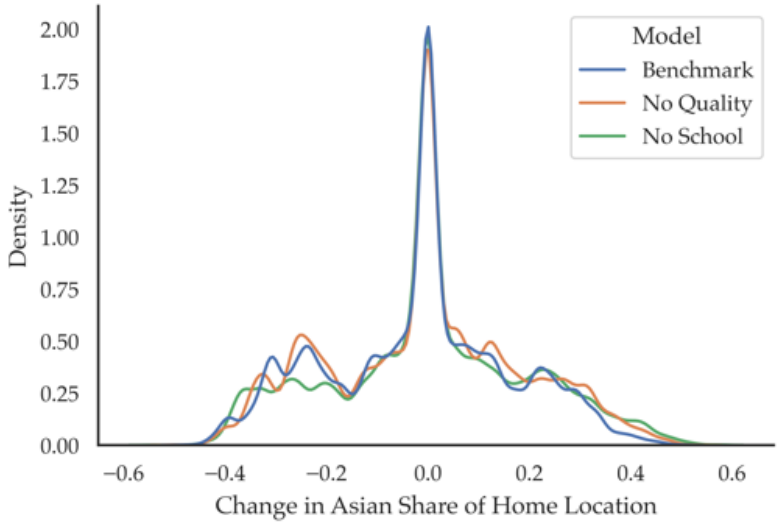

Note: See notes under Figure 9. This figure shows the equivalent statistic for Asian families who move and the change in the Asian share of their home locations.