

School Access and City Structure*

Giorgio Pietrabissa

LSE

October 14, 2026

Abstract

The location choices of households with children connect housing, labor, and school markets. I study how school catchment areas affect location choices, welfare, and city structure. I estimate a quantitative spatial model of the three markets for Madrid, and use a decomposition of market access to show that the policy binds most in the center, where schools are densest. Compared with free choice, catchment areas shorten school trips by about 40% but lower parents' welfare by about 7.5%, via lost schooling options. Parents suburbanize, and the city becomes more segregated and marginally more productive, with gains falling short of parents' losses.

*I am particularly grateful to Milena Almagro, Paula Bustos and Diego Puga for their help and guidance in the development of this project. This research was part of my PhD dissertation. I wish to thank the Consejería de Educación, Ciencia y Universidades of the Community of Madrid for access to the education data. I am also indebted to Gerardo Azor Martínez, Samuel Bentolila, Ignacio Berasategui, Caterina Calsamiglia and José Montalbán for help with the data. This paper benefited greatly from conversations with Gabriel Ahlfeldt, Pello Aspuru, Sara Bagagli, Tomás Budí-Ors, Felipe Carozzi, Francesco Chiochio, Jorge De la Roca, Cauê Dobbin, Nezihi Guner, Jessie Handbury, Vernon Henderson, Hans Koster, Eduardo Morales, Peter Morrow, Dávid Nagy, Henry Overman, Josep Pijoan-Mas, Milan Quentel, Andrés Rodríguez-Clare, Clara Santamaria, Marta Santamaria, Matteo Sartori, Olmo Silva, Jan Stuhler, Nick Tsivanidis, Jos Van Ommeren, Elisabet Viladecans-Marsal, Tom Zohar and Yanos Zylberberg, and seminar discussants and participants at CEMFI, UC Berkeley, Humboldt University, the European Meeting of the UEA, the UEA Summer School, the fRDB Fellows Workshop, the Junior Spatial Economic Conference, and the workshop on Public Policies, Cities and Regions at ENS Lyon. I acknowledge financial support from the María de Maeztu Unit of Excellence MDM-2016-0684, funded by MCIN/AEI/10.13039/501100011033. All errors are my own. E-mail: g.pietrabissa@lse.ac.uk

1 Introduction

Schools are an essential amenity that cities offer to families with children. Yet cities differ widely in how their policies determine access to schools. The differences reflect a tension between assigning students to schools near their homes and widening parental choice to schools farther away. Many cities opt for assignment through catchment areas (NCEE, 2026), often with the stated aim of shortening school trips and reducing the administrative burden of allocating seats. By tying school admission to residential location, however, catchment areas raise house prices near good schools and reinforce segregation across both neighborhoods and schools (Bayer et al., 2007; Black, 1999; Caetano and Macartney, 2021).

School admission policies, however, may matter well beyond schools and shape the structure of the whole city. They decide which schools are within reach of each home, and so steer the residential choices of households with children, potentially at the cost of longer commutes to work for parents. The consequences could also spill over to households without children, who compete for the same housing and jobs. Evaluating admission policies therefore requires tracing their effects from schools to the rest of the city. Still, no framework traces how admission policies propagate through the housing and labor markets in general equilibrium.

This paper studies school access and how introducing catchment areas for school admission affects welfare, as well as residential choices and labor market outcomes throughout the city. I estimate a quantitative spatial model for Madrid, where parents can apply to any publicly funded school. I introduce a new decomposition of market access into two margins and find that, under the estimated parameters, the spatial variation of school access stems from the number of schools within reach, rather than from their quality. I then compare the free-choice equilibrium with one in which catchment areas restrict the pool of schools available to parents. Catchment areas shorten school trips, but lower parents' welfare, lengthen their commutes, and make the city more segregated, while the gains in productivity and non-parents' welfare are modest relative to parents' losses.

Four facts from Madrid's travel and residential data suggest that accessing schools influences where parents live and how far they commute to work. First, households with children live farther from the city center than those without, especially among the high-skilled. Second, school trips account for 22% of all trips, behind only work, yet cover a median of only 2 kilometers against 10 for work. Third, students from high-skill households live closer to the school they attend than their low-skill peers, while high- and low-skill workers commute alike once household traits are controlled for. Finally, comparing workers at the same workplace, parents commute about 3% longer because of where they reside. Together, the facts motivate a model in which households differ by parenthood and skill, school trips are costly, and the residential choices of parents form a bridge between the school and labor markets.

To quantify how school access spills over to the rest of the city, I develop a quantitative spatial model of residence, workplace and school choice. The model builds on Ahlfeldt et al. (2015) and lets four household types, defined by the combination of parenthood and skill, choose where to live and work, and, if they have children, where to send them to school. School choice depends on quality net of travel costs and

on admission probabilities, which together define the school access each neighborhood offers. Quality responds to the skill mix of students through peer effects, so residential sorting and school quality feed back on each other. Firms compete with households for floorspace and combine it with high- and low-skill labor, benefiting from agglomeration. In equilibrium, house prices, wages, and school quality are jointly determined across space, and school allocation policies enter the model only through admission probabilities.

Understanding the impact of catchment areas first requires estimating the key ingredients of school access. These are the cost of traveling to school, the elasticity of school choice to quality, and the spatial distribution of schools and their quality. Their values set the potential bite of catchment areas, which would leave choices unchanged if, for example, parents only attended their nearest school even without constraints. Madrid's free choice system lets me recover all three ingredients from the universe of school admissions, where parents reveal how far they are willing to travel for a better school given their chances of admission. In the estimation results, school travel turns out to be about 2.5 times as costly as commuting to work, consistent with the shorter school trips in the data, yet low enough for parents to enjoy a varied pool of options. Parents also respond more strongly to differences in school quality than workers do to those in wages, with high-skill parents the most responsive.

The model estimation also necessitates exogenous sources of variation to separate residential preferences from peer effects in school quality. In the data, high-skill parents cluster around good schools either because they value them more or because their presence raises quality, and the two cannot be told apart directly. I break the feedback with two instruments drawn from Madrid's historical development, one for each direction. The first exploits how affluent households settled on higher ground to avoid the flow of wastewater before modern sewers were introduced. As a result, a school's elevation still predicts the skill mix of its students today through the surrounding demographics, while plausibly affecting quality only via the student composition.¹ The second instrument exploits the fact that newer school buildings are associated with higher perceived quality, and traces the original construction year of schools around a residential location. Conditional on other neighborhood characteristics correlated with the age of the building stock, the instrument shifts school access independently of sorting, and thus isolates residential preferences. Peer effects turn out small, with an elasticity of quality to the student skill mix just below 0.04, consistent with Madrid's equalized school funding, while the elasticity of residential choices to a neighborhood's utility falls within the range of previous estimates (Ahlfeldt et al., 2015; Miyauchi et al., 2025; Tsivanidis, 2026).

Perhaps counterintuitively, estimated school access is highest in the city center even though school quality is not, and a simple decomposition of market access explains why. Rationalizing the result calls for separating how many options a household can reach from how good they are, two dimensions that market access bundles together. I show that any market access measure in quantitative spatial models splits exactly into an extensive margin, the effective number of options (Adelman, 1969), and an intensive margin,

¹Historical sorting away from pollution explains persistent neighborhood sorting in Heblich et al. (2021), and Loumeau (2023) similarly instruments peer effects in Paris schools with longitude, exploiting the historical sorting of richer households away from the polluted East.

their average utility. The relative importance of each margin rests on the three ingredients of access: travel costs, the elasticity of choices to utility, and the spatial distribution of value across options. Their estimates for Madrid reveal that costly school trips, to which parents respond strongly, and scattered quality make school access peak in the center entirely through the extensive margin, i.e., the density of available schools. Workplace access peaks there too, but through the intensive margin, i.e., the average wage net of commuting costs, as commutes are cheaper and wages fall steeply away from the center. The decomposition is especially useful for the spatial analysis of goods not allocated through prices, where the allocation rule becomes a fourth ingredient that shifts the balance between the margins. Given the estimated ingredients, catchment areas in Madrid should bind most in the center, where the extensive margin is largest, and the common conditions behind this pattern suggest that it extends to other cities.

Since the spatial distribution of school access depends on the number of available options, I run a counterfactual imposing catchment areas, which directly limit the options available to parents, and then trace the citywide effects. The policy restricts enrollment to schools within the *barrio* of residence, an existing administrative unit of Madrid. Because the model abstracts from school capacity and lets floorspace adjust, I interpret the exercise as a comparison of long-run city equilibria under free choice and under catchment areas. I also vary the strength of peer effects and hold fixed, in turn, floorspace supply and the location of non-parents, to isolate the key general equilibrium mechanisms. Catchment areas along *barrios* cut school access by about 10% on average, but more in the center, so parents move toward the periphery, while non-parents react by moving in the opposite direction and partially offset the shift in residential demand. As intended by the policy, school trips shorten sharply, by about 40%.

Despite shorter school trips, parents' welfare drops under catchment areas, and both neighborhoods and schools become more segregated. Welfare falls by 7.4% for high-skill parents and by 7.6% for low-skill parents, while non-parents gain only about 0.2%. The losses stem from the restricted choice of schools and change little across scenarios, since the extensive margin behind them is directly constrained by the policy and is largely insensitive to general equilibrium channels. As high-skill parents sort more strongly into the areas with the best access and then into the best schools, residential segregation by skill among parents grows by 8.6%, and school segregation by 80%. Segregation across neighborhoods by parenthood rises even more, by 146%, as parents and non-parents relocate in opposite directions, and the relocation of non-parents alone accounts for about a third of the increase. Segregation increases along all three dimensions even when school quality is held fixed, and it climbs further as peer effects strengthen, documenting that the feedback between residential sorting and school quality magnifies the segregating force of catchment areas.

As residents relocate and floorspace adjusts, the policy spills over to the structure of the whole city, which becomes more suburbanized, more specialized, and ultimately more productive. Following residential demand, total floorspace densifies away from the center and diverges in its use, as suburbs turn more residential and the more productive central locations more commercial. The resulting agglomeration gains raise aggregate productivity by 0.07% and output by 0.26%, but they are modest relative to parents' welfare losses, which already account for the concurrent wage changes. These gains rest entirely on the endoge-

nous supply of floorspace through construction and its allocation across uses, and turn into a slight loss when both are held fixed. The endogenous response of non-parents also turns out to be a relevant mechanism, as they dampen the effects of the policy shock. Without them, for example, the price of floorspace would have risen more in the outskirts and dropped further in the center, resulting in a city-center gradient change about 50% larger. Finally, parents bear the cost of their relocation through a harsher trade-off between school and work trips, and their commutes rise by about 2 to 3%.

The results highlight an unfavorable trade-off between proximity and choice in school admission and thus call for a reassessment of the widespread use of catchment areas. They successfully shorten school trips, as intended by the policy, but evaluating them on trips alone hides their costs for the families they target, whose welfare falls, and for the city as a whole, which becomes more segregated. The accompanying gains in output and for non-parents are small relative to these losses. The trade-off worsens where school quality depends more on local demographics, as the losses of low-skill parents grow while those of high-skill parents shrink, and segregation rises sharply. It likely worsens in the short run as well, since floorspace cannot adjust and deliver the production gains, non-parents do not dampen the shock by relocating, and school capacity binds. Overall, the results confirm that school admission policies must be evaluated at the scale of the whole city to fully quantify their consequences.

Contribution and related literature. First, this paper contributes to the literature on school catchment areas by extending their study from local housing market effects to the whole city, including the labor market and households without children. A large literature shows that, by tying admission to residence, catchment areas capitalize school quality into house prices and carry residential sorting into the composition of schools, and that reforms of admission rules move prices, student outcomes, and residential choices accordingly (Black, 1999; Bayer et al., 2007; Gibbons et al., 2013; Caetano and Macartney, 2021; Billings et al., 2014; Machin and Salvanes, 2016; Campos and Kearns, 2024; Bibler and Billings, 2020).² These studies identify local effects, holding fixed what happens elsewhere in the city, from wages and commutes to the location choices of other households. Studies of school choice, in turn, largely take residential location as given (Abdulkadiroğlu et al., 2017; Agarwal and Somaini, 2018; Kapor et al., 2020; Laverde, 2020). Both approaches thus abstract from the citywide general equilibrium effects of catchment areas.

The equilibrium models closest to this paper study school allocation policies in settings with joint residential and school choices, but none features a labor market. Greaves and Turon (2024) compare geographic and lottery admissions in a stylized model of two adjacent neighborhoods, and show that households without children dampen the price premium around popular schools, a role non-parents also play here. Their model follows households over the life cycle, accounting for moving costs and for future parents who settle near good schools before having children, something the cross-sectional data of this paper do not allow. Agostinelli et al. (2024) estimate a spatial equilibrium model close to the one in this paper, but richer on the

²On capitalization, boundary designs compare homes on opposite sides of catchment borders (Black, 1999; Gibbons et al., 2013), and equilibrium sorting models recover the underlying preferences (Epple and Sieg, 1999; Sieg et al., 2004; Bayer et al., 2007); see Black and Machin (2011) and Kuminoff et al. (2013) for reviews. On reforms of admission rules, see Bogart and Cromwell (2000), Reback (2005), Ries and Somerville (2010), Lee (2015), and Machin and Salvanes (2016) for property prices, Lavy (2010) and Campos and Kearns (2024) for student outcomes, and Gortázar et al. (2023) for travel to school in Madrid. The link between schools and residential sorting has deep roots in public finance (Tiebout, 1956; Benabou, 1993; Fernandez and Rogerson, 1998).

education side, with children’s skills, private schools, and busing all entering school choice. This richness lets them evaluate both place-based and people-based policies, and they find that policies granting low-income families access to high-quality schools lower aggregate welfare, while distance and zoning shield the schools of affluent neighborhoods. [Hahm and Park \(2022\)](#) show that, under New York City’s centralized assignment, endogenous location choices halve the equity gains from removing geography-based priorities, but hold house prices fixed. As a result, none of these models can trace how catchment areas reach parents’ commutes, the floorspace that firms and households compete for, and thus the structure of the whole city.³ Doing so reveals effects beyond schools and housing: catchment areas make the city more suburbanized, more specialized, and slightly more productive, while marginally improving non-parents’ welfare.

Second, this paper extends quantitative models of city structure to school choice under admission rules that ration access, and to sorting by skill and parenthood. In the canonical models, access to jobs shapes where people live ([Ahlfeldt et al., 2015](#); [Heblich et al., 2020](#); [Tsivanidis, 2026](#)), and recent work adds access to consumption ([Miyauchi et al., 2025](#)). Extensions with heterogeneous households show how sorting and endogenous amenities shape the distributional effects of urban change ([Diamond, 2016](#); [Almagro and Domínguez-lino, 2019](#); [Couture et al., 2023](#)), while parenthood shapes location choices over the life cycle ([Moreno-Maldonado and Santamaria, 2021](#); [Albouy and Faberman, 2025](#)). Closest to this paper, [Loumeau \(2023\)](#) embeds school choice in a model à la [Ahlfeldt et al. \(2015\)](#), with housing and labor markets and peer effects in school quality, to evaluate where to locate new schools in the Paris region, and finds that a more decentralized strategy would raise welfare. In his framework, however, access is never rationed by admission rules, and households do not differ in skill. This paper brings admission rules into this class of models through admission probabilities estimated from the universe of school admissions in Madrid, and uses them to quantify the impact of catchment areas. Moreover, I let parents and non-parents of different skill sort across the city, tracing their welfare changes, and allow productivity to respond via agglomeration.

Finally, I introduce a decomposition of any market access measure into extensive and intensive margins. Market access is an ex-ante statistic capturing how households value the options they can reach, from jobs ([Ahlfeldt et al., 2015](#); [Tsivanidis, 2026](#)) to consumption ([Miyauchi et al., 2025](#)), but existing measures do not separate how many options a household can reach from how good they are. The decomposition builds on two ideas from other fields. The first is the inverse Herfindahl index from the IO tradition, which measures the number of equally important options ([Adelman, 1969](#)), and whose entropy-based counterpart [Caldwell and Danieli \(2024\)](#) use to measure workers’ outside options. The second is the split, drawn from trade and spatial economics, of a price index into the prices of available goods and their availability ([Feenstra, 1994](#); [Handbury and Weinstein, 2015](#)). Accordingly, I show that any market access measure splits exactly into the effective number of options and their average utility. The decomposition explains why school and

³[Eckert and Kleineberg \(2021\)](#) and [Fogli et al. \(2026\)](#) link location, education, and labor markets in models of intergenerational mobility, across US counties and between two neighborhoods of a representative metro area, respectively. In both, education policies reach labor markets through children’s future earnings, while here catchment areas reach them through parents’ current residential and commuting choices within the city.

workplace access are distributed so differently across the city. In the spirit of Campos and Kearns (2024), who find that the option value of an expanded choice set predicts where a reform improves educational outcomes most, I also use the decomposition to predict where catchment areas bind most.

The remainder of the paper is organized as follows. Section 2 describes Madrid’s school admission system and the data, and Section 3 documents how location and travel patterns differ by parenthood and skill. Section 4 develops the quantitative urban model with housing, labor, and school markets, and Section 5 estimates it. Section 6 decomposes school and workplace access into their margins, Section 7 quantifies the effects of catchment areas, and Section 8 concludes.

2 Context and data

2.1 The school system in Madrid

Schools in Spain fall into three types: public, concerted, and private. Public and concerted schools are both publicly funded and required to follow the same curriculum and admission rules. Concerted schools differ only in that they are privately managed⁴ and legally barred from charging fees for compulsory education.⁵ Private schools, by contrast, receive no public funding: they have more freedom to design their own curriculum, set their own admission criteria, and charge families directly.

Two features of Madrid’s system make it a suitable context for school choice estimation. First, families can apply to any public or concerted school in the city. Admission is governed by a transparent points system rather than by strict residential catchment areas. Parents can therefore consider a vast set of schools, letting me estimate school choice over space and determine the key ingredients of school access. More precisely, admissions depend solely on priority points, resolved through the Boston (or immediate acceptance) mechanism. The priority criteria include proximity to residence, siblings, and family income, applied uniformly across applicants, with no hard spatial constraint on choices.⁶ Parents can rank at most five schools in their applications, though they choose that set freely. In more than 90% of cases, a student is admitted to her most preferred option.

Once admitted, a student keeps her seat automatically in every subsequent year, without having to reapply or risk losing her spot. Competition through the points system hence concentrates almost entirely at entry. Around 70% of all school admissions in a given year come from students entering the first year of preschool, with the remaining 30% spread across the rest of preschool, primary, and secondary school.⁷

⁴Most concerted schools are administered by institutions linked to the Catholic Church, historically the dominant provider under the concerted system. In the municipality of Madrid, 70% of concerted schools are religious.

⁵In practice, most concerted schools collect payments framed as voluntary contributions for complementary services, and surveys find that families frequently report these payments are not truly optional. Because school quality is inferred from parents’ revealed choices, I account for any such unobserved price by allowing it to enter the quality directly.

⁶See Figure A.1 for the full list of criteria and points in years 2012–2015. Because acceptance depends on these priority points rather than stated preferences alone, families have an incentive to rank schools strategically, which also helps me simplify the school choice estimation strategy. The proximity-based points are scored when applying within the boundaries of a *distrito*, a large administrative unit that divides the municipality of Madrid into only 21 zones.

⁷See Figure A.2, which depicts the share for each compulsory school year between 2012 and 2015.

Madrid therefore differs substantially from settings where catchment areas tie each address to a single or very few assigned schools.

The second relevant feature of the system is that funding and curriculum are equalized across public and concerted schools by the regional government, narrowing their role in school choice. School quality will thus depend on other features, such as peer effects, teacher quality, complementary services, and the physical infrastructure. This separates Madrid from the US and other contexts, where school funding is often tied to local property taxes, linking school resources directly to the wealth of the surrounding neighborhood.

2.2 Data

I assemble a dataset from several administrative and survey sources, spanning residential registries, school admissions records, and travel surveys. Together, they let me trace how households in Madrid choose where to live, where to work, and where to send their children to school, by their educational level and parenthood status. Throughout the paper, I use the *sección censal* administrative area as the geographical unit, and note explicitly whenever an analysis relies on a different one.⁸

School choices. Administrative data from the Community of Madrid on school applications and admissions cover the universe of students applying to public and concerted schools. Each record includes the residential location, the first-choice school, the school of admission, the application score, and other characteristics relevant to admission. The available data cover 2012 to 2015. The dataset does not report household educational level, so I proxy it using the share of households with at least one child aged two or younger that hold a college degree or higher, computed within the family's *sección* of residence from the residential administrative data described next.

Residential choices. Household counts by group come from the Padrón Municipal of the city of Madrid, the municipal residency registry that every resident must join to access public services. Each individual reports nationality, place and date of birth, sex, educational attainment, and address, and cohabitants are linked together regardless of whether they are related. The data cover 2010 to 2018 and are available at the *sección* level. I define a family as a household with at least one person aged 17 or younger, and a high-skill household as one with at least one member holding a university degree or higher.

Workplace choices. Commuting information comes from the 2018 Encuesta Domiciliaria de Movilidad (EDM2018), a mobility survey run by the Transport Consortium of the Community of Madrid. Respondents report all trips taken on a given day, including trip purpose, travel time, origin, destination, and demographic characteristics such as educational attainment. I construct skill-specific commuting matrices across *barrios*, and include municipalities in the wider urban area for the estimation of commuting elastic-

⁸ *Secciones* are Spain's smallest administrative division, and the average *sección* in Madrid had just over 520 households in 2017. Where a coarser unit is needed, I aggregate to *barrios*, the next level up. Above these are the *distritos*. I adopt the geographical boundaries of 2017, under which Madrid's 2420 *secciones* group into 128 *barrios* (about 19 *secciones* each) and 21 *distritos*. *Distritos* are used for proximity points in the current enrollment system (see Figure A.1). For the analysis of commuting, I also include the other 51 municipalities of the Madrid metropolitan area.

ities. To get the distribution of workers across workplaces, I use administrative data from Social Security records, available at the *barrio* level. To obtain the finer distribution at the *sección* level, I split each *barrio* total across *secciones* in proportion to commercial floorspace from the cadastre.

Travel times. I construct travel times for each origin-destination pair across Madrid. For *barrios* and municipalities in the metropolitan area, I query Google Maps for walking, driving, and transit times at 8:30 a.m. on a weekday. For the matrix between *secciones*, I use the Open Source Routing Machine for walking and driving, and approximate transit times with those of the corresponding *barrio*-level pair. I then take a weighted average across travel times using predicted mode shares for school and work trips separately, further conditioned on skill and distance, as estimated from the EDM2018 travel survey. See Appendix B.1 for more details on the construction of travel times.

Other data. I complement the main sources with additional data, mostly at the *sección* level. Cadastral records separate built structures into residential and commercial use. Residential rental prices come from tax declarations compiled by the national statistical office (INE), and local income, which comes from the same tax records. For commercial structures I rely on cadastral values, available only at the *barrio* level. Household expenditure shares by type come from the Encuesta de Presupuestos Familiares (EPF), the consumption survey by INE. Finally, I obtain school-level test results and vacancy counts from the Community of Madrid, and complement them with information on schools' curricula and services scraped online.

3 Motivating evidence: location and travel patterns by parenthood and skill levels


This section describes how parenthood and skill levels shape the location and travel choices of households in Madrid. I first show that parenthood sharply separates residential patterns within the city. I then link the parenthood divide to school trips, which fall almost entirely on households with children and are highly concentrated in space. Finally, I find that high-skill families sort more strongly toward schools, while parenthood pushes households to live farther from their jobs. The facts together show that households sort within the city according to their parenthood status and skill level, and that the need to access schools likely shapes their spatial choices, including where they live and work.

To build the motivating evidence in this section, I rely on two data sources. First, the Padrón Municipal, for residential counts of households by group, and second the EDM2018, for trips by household type and destination. See Section 2.2 above and Appendix B for more details on the data and its construction.

Fact 1. The city-center residential gradient differs by parenthood and skill level.

In Madrid, households without children are more likely to live toward the city center, while the opposite holds for households with children. Taking the cross-section of household counts across *secciones censales* in 2019 and ordering them by distance to the center reveals the pattern clearly. Figure 1 plots the city-center gradient with a polynomial prediction for each household type. Only high-skill households without children display a markedly negative gradient, with a correlation of -0.35 , while all other types show a positive relationship. Most importantly, parenthood pushes households farther from the center, for both

Figure 1: Residential choices and distance to the center.



Notes: The graphs show the distribution of residential shares for each household group against the distance to the center in a smooth polynomial. I also plot the scatter dots for each *sección censal*. Panel (a) includes high-skill households, while Panel (b) reports low-skill households. I identify the center of Madrid with *Puerta del Sol*, and cut the data at 14 km as for longer distances there are too few observations for the polynomial prediction (it drops 46 units out of 2420). The graphs include the raw correlations between the residential shares and distance to the center for exposition.

high- and low-skill households, and the shift is considerably more pronounced for the high-skilled.

Many forces could drive this sorting: the availability of different housing types, the location of amenities catering to each group, or different sensitivities to house prices. This paper focuses on school accessibility, which I next show matters primarily for the mobility of households with children.


Fact 2. Trips to school are quantitatively relevant and spatially concentrated.

Mobility patterns reveal that school trips are a first-order travel reason. In Panel a of Figure 2, I use the EDM survey to plot the share of trips by motive. For the city as a whole, commuting to work is the most frequent reason to move in Madrid, at 29% of all trips. Traveling to school is a close second at 22%, with the remaining half split among categories that mostly capture leisure and consumption. The aggregate share, however, masks a sharp divide: school trips fall almost entirely on households with children, the only group whose spatial choices must account for the location of schools.

School trips are also markedly shorter than work commutes. In Panel b of Figure 2, I compare the two travel-time distributions.⁹ The school distribution is skewed toward lower values: the median school trip takes around 12 minutes, against more than 30 for work. The same pattern holds for travel distances, so the gap is not merely the result of mode choice and travel speeds, but reflects the greater spatial concentration of school trips (see Figure A.3). Households with children therefore do not travel far to reach schools, despite the absence of restrictions on school access in Madrid. Several forces could explain the pattern: their residential sorting, a distaste for traveling to school, or a homogeneous and evenly distributed supply of schools across the city. I next show that sorting towards schools may be relevant.

⁹Apart from within-unit pairs (Appendix B.1), I avoid the travel times reported in the EDM survey, which bunch systematically at multiples of 5 minutes, as shown in Figure A.3. Instead, I use travel times computed from OSRM and Google Maps, combined with the predicted mode shares from the EDM by type of trip and skill group. For more details, refer to the Data section.

Figure 2: Trips by motive: relevance and times



Notes: In Panel (a), the graph shows the share of each type of trip by their motive. It uses trip-level data from the EDM. The first breakdown includes all households, reflecting the overall share of trips in an average day in the metropolitan area Madrid. The second breakdown conditions only on households with children, and the third on those without. Work and school trips are highlighted in colors, the others are shaded in gray for visual purpose. The motive definitions are determined by the EDM. Return trips to the residential address are excluded. In Panel (b), the graph reports the separate distributions of travel times for trips to a workplace or a school. The dashed lines are set at the median value of each distribution. Travel times are computed from Google Maps, using the origin and destination *barrios*, and averaging car, transit and walking times with mode weights predicted from the EDM survey, separately by motive.

Fact 3. Skill levels matter for sorting toward schools, but not toward workplaces.


Next, I show that skill drives sorting toward schools but not toward jobs. Using the trips data from the EDM survey, I estimate the effect of belonging to a high-skill household on travel times:

$$\ln \text{Travel Time}_i = \beta \text{High-skill HH}_i + \gamma X_i + \epsilon_i. \quad (1)$$

Figure 3 reports β across four specifications with progressively richer controls. Focusing on trips to school (in red), students from high-skill households travel 10% longer to school than their low-skill counterparts, a gap that halves to 5% once household characteristics are added. Conditioning on the residential neighborhood, by including a home fixed effect, erases the difference: from the same location, both groups reach equally distant schools. Conditioning instead on the location of the attended school, by including a destination fixed effect, flips the sign of the coefficient: students from high-skill households live closer to their school than their low-skill counterparts. High-skill households thus travel longer to school on average, yet sort more strongly toward the schools they attend.

The same specification on travel times to work (in green), now including households with and without children, reveals no comparable sorting toward workplaces. Workers from high-skill households take 5% longer to reach their jobs, but the difference fades once demographic controls are added and stays at zero after conditioning on residential or workplace fixed effects. The contrast with school trips reflects the difference in average travel time documented in Fact 2. It points towards schools as a potentially stronger force for sorting, as highlighted by high-skilled households responding more strongly than low-skilled ones.

Figure 3: The effect of skills and parenthood on travel times.



Notes: The graph reports the coefficients of four specifications for each of three empirical models, with 95% confidence intervals. The first group (red) gives the effect of belonging to a high-skill household on travel times to school; the second (green), the same effect on travel times to work; the third (yellow), the effect of belonging to a household with children on travel times to work. The unit of observation is the individual trip from the EDM survey. Travel times are measured via Google Maps between origin and destination *barrios*, averaging car, transit, and walking times with predicted weights by travel motive and skill. The first specification (circle) is the raw effect without controls. The second (square) adds controls: age group, sex, education level, Spanish nationality, reduced mobility, household size, number of vehicles, carers, and workers, mode of transport, and, for workers only, job activity, sector, and occupation. The third (triangle) adds a home *barrio* fixed effect; the last (diamond) replaces it with a destination *barrio* fixed effect. Table A.1 in the Appendix reports the regression results.

Fact 4. Households with children sort away from workplaces.

Finally, I show that having children lengthens work commutes, through residential sorting. The last group of coefficients in Figure 3 (in yellow) reports the specification capturing the effect of belonging to a household with children on travel time to work,

$$\ln \text{Travel Time}_i = \beta \text{HH with children}_i + \gamma X_i + \epsilon_i. \quad (2)$$

The unconditional gap is around 6%, but accounting for household characteristics already renders it indistinguishable from zero. The same occurs when including the home fixed effect. Instead, including the destination fixed effect, and thus comparing two commuters who work in the same location, the difference turns positive again. The parenthood commuting gap therefore reflects, beyond household characteristics, a difference in residential choices. The pattern echoes the previous facts: households with children sort differently, likely pulled by the need to reach schools, a force made powerful by how important and how short school trips are.

In conclusion, these facts open the door to substantial general equilibrium effects of school policy on the

city, through a chain of spatial choices. First, schools are a major, and typically short-range, destination for the mobility of households with children. Second, and for these reasons, they likely draw this group's residential choices in a different direction from the rest of households, with a strength that varies by skill. Third, the sorting may reach beyond the housing market through the commuting response of parents, and spill over to the labor market as well.

The full chain, however, is only suggested by the facts above. To study its relevance in a unified framework, I next develop a quantitative urban equilibrium model in which households choose their residence, their workplace, and, if they have children, their school. Once estimated, the model quantifies the strength and far-reaching effects of school access policies on city structure.

4 A quantitative urban model with housing, labor and school markets

In this section, I develop a general equilibrium quantitative urban model of the residential choices of households who differ in skills and parenthood. The model features three modules: a housing market in which households make residential choices across locations, a labor market in which they choose workplaces, and a school market in which parents choose where to enroll their children. Given its primitives, the model characterizes the distribution of housing prices, wages, and school quality across locations. These objects are determined in general equilibrium, since they both affect and respond to the spatial distribution of households in the city.

The model closely follows the seminal work of [Ahlfeldt et al. \(2015\)](#). It departs from that framework mainly by allowing for different household types, as in [Tsivanidis \(2026\)](#), and by adding a school market, which introduces an additional travel choice, as in [Miyauchi et al. \(2025\)](#).

4.1 Set up

The economy is a closed city with a fixed number of neighborhoods $n = 1, \dots, N$, each with a fixed amount of land. A fixed measure of households $\bar{R}$ populates the city, ex-ante heterogeneous along two dimensions: education level, high or low ($g \in \{L; H\}$), and family status ($k \in \{0; 1\}$), indicating whether the household has children. Mobility across neighborhoods depends on geography and transport technologies, reflected in exogenous travel times.

4.2 Preferences

A household ω living in neighborhood n chooses consumption of a freely traded numeraire (C_n) and of residential floorspace (H_n) to maximize utility, with Cobb-Douglas preferences that vary across types g, k . Floorspace carries a price in each neighborhood (r_n), and rents collected across the city are redistributed equally to households through an endogenous lump-sum transfer π .¹⁰ All households value local ameni-

¹⁰Landlords are absentee and do not consume floorspace themselves. Rents collected across the city are fully redistributed to households as a uniform lump-sum transfer, $\pi \equiv \frac{\sum_n r_n H_n^R + \sum_i r_i^C H_i^C}{\bar{R}}$, rather than being lost.

ties in their residential neighborhood, B_{ngk} . Households work in some location $i = 1, \dots, N$, bearing a commuting cost between residence and workplace (τ_{nig}^W). They earn wages w_{ig} that depend on skill group g and vary across workplaces i . If the household has children ($k = 1$), utility also depends on the quality of the school its children attend (S_{jg}), which varies across skill groups g and individual schools j . Reaching a school requires paying a travel cost from the residence (τ_{njg}^S). If the household has no children ($k = 0$), it does not value schooling in the utility maximization problem.

Every household is subject to an idiosyncratic shock for each location choice it makes: $b_{ng}(\omega)$ for the residence, $z_{ig}(\omega)$ for the workplace, and $s_{jg}(\omega)$ for the school. Each shock is drawn from a Fréchet distribution that varies across skill groups. Then, the indirect utility of household ω choosing residence n , workplace i , and school j is

$$V_{nijgk}(\omega) = \frac{B_{ngk} b_{ng}(\omega)}{r_n^{1-\alpha_{gk}}} \left(\frac{w_{ig} z_{ig}(\omega)}{\tau_{nig}^W} + \pi \right) \left(\frac{S_{jg} s_{jg}(\omega)}{\tau_{njg}^S} \right)^k, \quad k \in \{0; 1\}, \quad g \in \{H; L\}. \quad (3)$$

Travel costs are exponential in travel time, $\tau_{nig}^W = e^{\kappa^W t_{nig}}$ and $\tau_{njg}^S = e^{\kappa^S t_{njg}}$, where t denotes travel time and κ , common across skill groups, is the semi-elasticity of travel costs to travel time. Since choice probabilities depend on $\tau^{-\theta_g} = e^{-\theta_g \kappa t}$, the semi-elasticity of choice probabilities to travel time is $\phi_g^W \equiv \theta_g^W \kappa^W$ for workplaces and $\phi_g^S \equiv \theta_g^S \kappa^S$ for schools.

This utility structure rests on four simplifying assumptions. First, the model abstracts from dynamics in location choices, and in particular from those related to family formation. Second, the model abstracts from household internal structure, and each household behaves as a single agent. Third, workplace and school choices are independent, with no scope for chain trips across the city. Fourth, households first choose where to live, and then where to work and where to send their children to school. This timing allows solving the location choice problem by backward induction: when choosing a residential neighborhood, households already account for access to workplaces and to schools.¹¹

4.3 Workplace choice and labor supply

The household then chooses where to work, independently of its other location choices. I assume wages, commuting costs, and workplace preferences do not vary by parenthood, so labor market choices are independent of k . The Fréchet idiosyncratic shock delivers bilateral commuting probabilities:

$$\lambda_{nig|n}^W \equiv \Pr \left[\frac{w_{ig} z_{ig}(\omega)}{\tau_{nig}^W} > \max_{i' \neq i} \frac{w_{i'g} z_{i'g}(\omega)}{\tau_{ni'g}^W} \right] = \frac{(w_{ig}/\tau_{nig}^W)^{\theta_g^W}}{\sum_l (w_{lg}/\tau_{nlg}^W)^{\theta_g^W}}, \quad \text{with } z_{ig}(\omega) \sim e^{(-z^{-\theta_g^W})}. \quad (4)$$

¹¹ All these assumptions are driven by data limitations, rather than expecting they capture household behavior better. I have no data on the residential choices of individual households over time; I have no information connecting the work or school choices of individual household members to the residential and school administrative data; I only observe school and work trips from separate data sources and cannot observe trip chaining; and I have no information on which choice is made first, so I take the choice-ordering assumption from the literature.

Taking the expectation over the idiosyncratic shock yields workplace access, the expected labor utility a household of skill type g obtains from residing in n :

$$\mathbf{W}_{ng} \equiv \mathbb{E}_i \left[\frac{w_{ig} z_{ig}(\omega)}{\tau_{nig}^W} \right] = \Gamma \left(1 - \frac{1}{\theta_g^W} \right) \left[\sum_i (w_{ig} / \tau_{nig}^W)^{\theta_g^W} \right]^{\frac{1}{\theta_g^W}} = \gamma_g^W (\Phi_{ng}^{RW})^{\frac{1}{\theta_g^W}}. \quad (5)$$

Aggregating commuting probabilities across residents gives the skill-specific labor supply to each workplace:

$$L_{ig} = \sum_n \frac{(w_{ig} / \tau_{nig}^W)^{\theta_g^W}}{\Phi_{ng}^{RW}} R_{ng} = w_{ig}^W \Phi_{ig}^{WR}. \quad (6)$$

Here, R_{ng} is the endogenous measure of type- g households living in n . Equation (6) relates workplace access for firms (Φ_{ig}^{WR}) to workplace access for workers (Φ_{ng}^{RW}), as in [Tsivanidis \(2026\)](#).

4.4 School choice and school demand

Parents ($k = 1$) choose a school for their children by trading off school quality against the travel cost of the daily trip. School assignment in Madrid, however, is centrally regulated, and the planner allocates seats according to priority rules, which introduce the probability of admission $p_{\omega j}$ as an additional object into the household's problem.

Following [Calsamiglia et al. \(2020\)](#), I assume each family submits a single school choice rather than a ranked list. If the application scores enough priority points, the student is admitted; otherwise, the student is assigned the preferred school among a pool of leftover schools. The school choice problem is:

$$\max_j \left\{ p_{\omega j} \frac{S_{jg} S_{jg}(\omega)}{\tau_{njg}^S} + (1 - p_{\omega j}) \max_{\ell \in \text{leftovers}} \left[\frac{S_{\ell g} S_{\ell g}(\omega)}{\tau_{n\ell g}^S} \right] \right\}. \quad (7)$$

Let $q_{\omega j}$ denote the priority points household ω obtains for school j , and $\bar{q}_j$ the minimum points required for admission. The probability of admission is an equilibrium object, determined by demand for the school, its fixed seat supply, and the resulting cutoff $\bar{q}_j$: $p_{\omega j} = \Pr\{q_{\omega j} \geq \bar{q}_j\}$. Assuming households have perfect information and are fully sophisticated, this probability becomes an indicator, as households know whether their own points score higher than the cutoffs. The problem then simplifies to choosing the preferred school within the set where admission is certain, $\mathcal{J}_\omega \equiv \{j | \mathbb{1}(q_{\omega j} \geq \bar{q}_j) = 1\}$. As with workplace choice, the Fréchet shocks deliver a closed-form logit expression for choice probabilities:

$$\lambda_{\omega jg|n}^S \equiv \Pr \left[\frac{p_{\omega j} S_{jg} S_{jg}(\omega)}{\tau_{njg}^S} > \max_{j' \neq j} \frac{p_{\omega j'} S_{j'g} S_{j'g}(\omega)}{\tau_{nj'g}^S} \right] = \frac{p_{\omega j} (S_{jg} / \tau_{njg}^S)^{\theta_g^S}}{\sum_\ell p_{\omega \ell} (S_{\ell g} / \tau_{n\ell g}^S)^{\theta_g^S}}, \text{ with } s_{jg}(\omega) \sim e^{(-s - \theta_g^S)}. \quad (8)$$

The admission probability $p_{\omega j}$ is the object through which the admission policy operates. It introduces a wedge in (8) that generates zeros for household-school pairs where admission is infeasible.

Taking the expectation over the idiosyncratic shock yields school access, the expected schooling utility

a household of skill type g obtains from residing in n . I assume households do not yet know their own priority points at the time of the residential choice. I therefore approximate the admission probability entering (9) as the share of households from n who have j in their certain admission set:¹²

$$\mathbf{S}_{ng} \equiv \mathbb{E}_j \left[\frac{p_{\omega j} S_{jg} s_{jg}(\omega)}{\tau_{njg}^S} \right] \approx \gamma_g^S \left[\sum_j \bar{p}_{nj} (S_{jg} / \tau_{njg}^S)^{\theta_g^S} \right]^{\frac{1}{\theta_g^S}}, \quad \text{with} \quad \bar{p}_{nj} = \frac{\sum_{\omega \in n} p_{\omega j}}{\sum_{\omega} \mathbb{1}\{\omega \in n\}}, \quad (9)$$

where the local admission probability $\bar{p}_{nj}$ is assumed not to vary with skill.¹³ School access nests the residential value of schooling under many assignment regimes, including catchment-area assignment, free school choice, and intermediate systems, through the admission probabilities that determine each household's relevant school set. Finally, aggregating across households gives the skill-specific demand, as a function of school quality:

$$\mathbf{K}_{jg} = \sum_{\omega \in g} \lambda_{\omega jg}^S = \sum_{\omega \in g} \frac{p_{\omega j} (S_{jg} / \tau_{njg}^S)^{\theta_g^S}}{\sum_{\ell} p_{\omega \ell} (S_{\ell g} / \tau_{n \ell g}^S)^{\theta_g^S}} = S_{jg}^{\theta_g^S} \sum_{\omega \in g} \frac{p_{\omega j} / (\tau_{njg}^S)^{\theta_g^S}}{\Phi_{\omega g}^S}. \quad (10)$$

4.5 Residential choice and demand

Having solved workplace and school choice, the household chooses where to live by maximizing expected indirect utility across neighborhoods. Substituting workplace access $\mathbf{W}_{ng}$ and school access $\mathbf{S}_{ng}$ from equations (5) and (9) into (3) gives the indirect utility of residing in n :

$$V_{ngk}(\omega) = \frac{B_{ngk} b_{ng}(\omega)}{r_n^{1-\alpha_{gk}}} (\mathbf{W}_{ng} + \pi) \mathbf{S}_{ng}^k. \quad (11)$$

The Fréchet assumption on $b_{ng}(\omega)$ delivers residential choice probabilities:

$$\lambda_{ngk} \equiv \Pr \left[V_{ngk}(\omega) > \max_{n' \neq n} V_{n'gk}(\omega) \right] = \frac{1}{\Phi_{gk}} \left(\frac{B_{ngk}}{r_n^{1-\alpha_{gk}}} (\mathbf{W}_{ng} + \pi) \mathbf{S}_{ng}^k \right)^{\theta_g^B}, \quad \text{with } b_{ng}(\omega) \sim e^{(-b^{-\theta_g^B})}. \quad (12)$$

Neighborhoods with better amenities, lower rents, and better access to work and school attract more residents. The measure of type- (g, k) households living in n follows directly, $R_{ngk} = \lambda_{ngk} \bar{R}_{gk}$. Aggregating residential stocks and housing expenditure across types gives demand for residential floorspace:

$$H_n^R = \sum_{g,k} (1 - \alpha_{gk}) \frac{\mathbf{W}_{ng} + \pi}{r_n} \lambda_{ngk} \bar{R}_{gk}. \quad (13)$$

¹²The approximation is required to aggregate school access at the household-type level, from ω to ng . Appendix D.3 investigates the approximation bias and shows it to be empirically negligible.

¹³Appendix D.2 discusses the assumption that admission probabilities are skill-independent, and provides empirical evidence in support.

Taking the expectation over the idiosyncratic shock gives the expected utility from living in the city for each type:

$$\bar{V}_{gk} \equiv \mathbb{E}_n[V_{ngk}(\omega)] = \Gamma\left(\frac{\theta_g^B - 1}{\theta_g^B}\right) \left[\sum_n \left(\frac{B_{ngk}}{r_n^{1-\alpha_{gk}}} (\mathbf{W}_{ng} + \pi) \mathbf{s}_{ng}^k \right)^{\frac{\theta_g^B}{\theta_g^B - 1}} \right]^{\frac{1}{\theta_g^B}} = \gamma_g^B (\Phi_{gk})^{\frac{1}{\theta_g^B}}. \quad (14)$$

Equations (12) and (14) are specific to both skill group g and family status k : workplace access varies by g through wages, school access enters only for parents ($k = 1$), and both the housing expenditure share α_{gk} and amenities B_{ngk} vary flexibly by type. Residential sorting along both dimensions follows in equilibrium.

4.6 Supply side: firms, floorspace and school quality

A representative firm in each workplace i produces the numeraire good using commercial floorspace and a CES bundle of high- and low-skill labor:

$$Y_i = A_i L_i^\beta (H_i^C)^{1-\beta}, \quad L_i = \left[\psi_i (L_{iH})^{\frac{\rho-1}{\rho}} + (1 - \psi_i) (L_{iL})^{\frac{\rho-1}{\rho}} \right]^{\frac{\rho}{\rho-1}}, \quad (15)$$

where ρ governs substitutability between skill types and ψ_i is a location-specific weight on high-skill labor. Profit maximization under perfect competition gives standard factor demands and a zero-profit condition for the price of commercial floorspace, r_i^C , derived in Appendix C. Productivity is subject to agglomeration economies:

$$A_i = \left(\frac{L_i}{T_i} \right)^{\delta^A} v_i^A, \quad (16)$$

where δ^A is the elasticity of productivity to density and v_i^A an exogenous productivity fundamental.

Floorspace is built by profit-maximizing developers who combine land with capital, supplied competitively in a single economy-wide market:

$$H_n = K_n^{\mu_n} (T_n \zeta_n)^{1-\mu_n}, \quad (17)$$

where ζ_n is a location-specific land productivity fundamental, and $\mu_n \in \{\mu_C, \mu_S\}$ is the elasticity of floorspace to capital, allowed to differ between the central city and the suburbs. Profit maximization yields a closed-form floorspace supply, increasing in the local price (Appendix C). Residential and commercial floorspace draw on the same underlying stock H_n , so the residential share σ_n is not chosen separately but adjusts to clear both markets. It equals residential expenditure relative to total floorspace expenditure, once commercial expenditure is converted to residential-equivalent units through a location-specific wedge ξ_n .¹⁴

$$\sigma_n = \frac{E_n^R}{E_n^R + E_n^C / \xi_n}, \quad (18)$$

where $E_n^R \equiv r_n H_n^R$ is residential floorspace expenditure from equation (13), and $E_n^C \equiv r_n^C H_n^C$ is commercial

¹⁴The wedge links the two rents, $r_n^C = \xi_n r_n$, and accounts for the observed gap between commercial and residential floorspace prices within the same location.

floorspace expenditure from equation (C.1). Workplaces and neighborhoods index the same set of locations, so i and n denote the same location here. Residential and commercial supply are then $H_n^R = \sigma_n H_n$ and $H_n^C = (1 - \sigma_n) H_n$.

School supply, in terms of number and locations of schools, is assumed fixed. Quality, however, adjusts according to peer effects:

$$S_{jg} = \left(\frac{K_{jH}}{K_{jL}} \right)^{\delta_g^S} v_{jg}^S, \quad (19)$$

where δ_g^S is the elasticity of quality to the student mix and v_{jg}^S is an exogenous quality fundamental for each school. This specification is agnostic on the micro-foundation of the determinants of school quality and thus captures any mechanism that works through the skill-mix of students' families.

4.7 Equilibrium

Given exogenous city fundamentals $\{v_i^A, \psi_i, v_{jg}^S, B_{ngk}, t_{nig}^W, t_{njg}^S, p_{\omega j}, \xi_n, T_n, \zeta_n\}$, model parameters $\{\alpha_{gk}, \beta, \rho, \mu_C, \mu_S, \delta^A, \delta_g^S, \theta_g^B, \theta_g^W, \theta_g^S, \kappa^W, \kappa^S\}$, and the aggregate stock of each household type $\bar{R}_{gk}$, an equilibrium is a set of wages $\{w_{ig}\}$, floorspace price $\{r_n\}$, and school quality $\{S_{jg}\}$ such that:

1. Labor markets clear: labor supply from equation (6) equals labor demand from equation (C.1), for every workplace i and skill group g .
2. Floorspace markets clear: given $r_n^C = \xi_n r_n$, residential demand from equation (13) equals residential supply $\sigma_n H_n$, and commercial demand from equation (C.1) equals commercial supply $(1 - \sigma_n) H_n$, in every neighborhood n .
3. The school market clears: enrollment from equation (10) is consistent with endogenous school quality (19) for every school j and skill group g .
4. The city is closed: $\sum_n \lambda_{ngk} = 1, \forall g, k$.

5 Model estimation

This section uses data from Madrid to estimate the model's parameters. The focus is on parameters governing household demand for residential locations, workplaces, and schools. Workplace and school choice parameters are obtained separately first, residential ones second, and wages, school quality, and residential amenities are recovered via revealed preference.

5.1 Workplace choice

Projecting commute shares on travel times separately by skill group recovers the semi-elasticity of commute choices to travel times from Equation 4. I estimate it by PPML (Santos Silva and Tenreyro, 2006), which handles the many zeros in the commuting matrices, and include origin and destination fixed effects.

Table 1: Commuting gravity equation

	PPML			Control-function IV	
	All (1)	High skill (2)	Low skill (3)	High skill (4)	Low skill (5)
Travel time	-0.085*** (0.003)	-0.076*** (0.003)	-0.093*** (0.002)	-0.082*** (0.002)	-0.096*** (0.002)
Origin FE	✓	✓	✓	✓	✓
Destination FE	✓	✓	✓	✓	✓
First-stage F				93785.1	115379.8
Observations	32,041	32,041	32,041	32,041	32,041

Notes: The dependent variable is the share of commuters from an origin to a given destination. Travel times are measured in minutes and differ between the high- and low-skill estimations, as each is a skill-specific weighted average across transport modes. Columns (1)-(3) report PPML estimates with two-way origin and destination clustered standard errors in parentheses. Columns (4)-(5) instrument travel time with straight-line distance through a control function, with standard errors bootstrapped over 100 replications. The geographic units are the *Barrios* and municipalities of the Urban Area of Madrid, for a total of 179 locations. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

I also instrument travel times with straight-line distance to purge the estimate of endogenous variation from congestion and from road and transit planning.¹⁵ Table 1 reports the results, which show low-skill workers are more sensitive to travel times, consistent with the literature (Tsivanidis, 2026; Lee and Tan, 2026). Robustness Table D.1 shows, among other checks, that splitting the sample by parenthood does not generate significantly different semi-elasticities.

With the estimated semi-elasticity, I use Equation 6 to recover the transformed wages $w_{ig}^{\theta_g^W}$ that rationalize the observed distribution of workers across *secciones*.¹⁶ Intuitively, locations that have more workers than predicted by the geography of the city and the distribution of residents must offer higher wages.

Next, to separate ϕ_g^W into its components, I impose a common κ^W across skill groups and match the variance of residential income to its model counterpart, the variance of workplace access aggregated across residents.¹⁷

$$\min_{\theta_H^W} \left\{ \text{Var}(I_n) - \text{Var}\left(\frac{R_{nH}}{R_n} W_{nH} + \frac{R_{nL}}{R_n} W_{nL}\right) \right\}^2. \quad (20)$$

Here, I_n is residential income in *sección* n normalized by its geometric mean, R_{ng}/R_n the share of skill- g residents, and W_{ng} their workplace access from Equation 5, which depends only on ϕ_g^W and θ_g^W . The

¹⁵Because PPML is non-linear, the IV strategy follows the control function approach with bootstrapped standard errors.

¹⁶The distribution of workers across *secciones* is built from Social Security employment counts by *barrio*, split into skill groups using the high-skill employment shares from the EDM2018, and allocated across the *secciones* of each *barrio* in proportion to their commercial floorspace.

¹⁷A common κ^W is a strong assumption, but less consequential given that travel times are already skill-specific through mode shares. Under a common κ^W , one can write $\theta_L^W = \theta_H^W \frac{\phi_L^W}{\phi_H^W}$. The variance match identifies θ_H^W ; κ^W then follows from ϕ_H^W , and θ_L^W from ϕ_L^W . Worker and resident counts are measured per unit of built floorspace, so that the moment reflects the spatial intensity of economic activity rather than differences in the physical size of *secciones*.

variance matching uses the dispersion of residential income to discipline that of transformed wages. The procedure yields $\kappa^W = 0.019$, while the dispersion parameters $\theta_H^W = 4.366$ and $\theta_L^W = 5.087$ are the PPML control-function IV estimates from the commuting gravity regression in Table 1. I then recover the wage vectors and compute workplace access from Equation 5.

5.2 School choice

The estimation of school preferences follows the same structure as workplace choice above. I use the admission data to construct the set of all student-school pairs in each year. I focus on students in the first year of pre-school, the point of entry into the school system and the only application most parents make over the school career of their children.¹⁸

I estimate the school-choice parameters by fitting observed choices with Equation 8 by maximum likelihood, following Bayer et al. (2007). Because I do not observe a family’s skill directly, I approximate it with the share of high-skill households with children aged two or younger who live in the same *sección* ($\hat{h}_\omega$).¹⁹ The likelihood therefore mixes the two skill types:

$$\mathcal{L} = \sum_{\omega} \sum_{j \in \mathcal{J}_{\omega}} \mathcal{I}_{\omega j} \ln \left(\hat{h}_{\omega} \lambda_{\omega j H}^S + (1 - \hat{h}_{\omega}) \lambda_{\omega j L}^S \right), \quad (21)$$

where $\mathcal{I}_{\omega j}$ indicates the school j that household ω actually chose.

The likelihood in Equation 21 is the individual-level counterpart of the aggregate gravity regression for workplace choice. Here, the student-specific admission indicator $p_{\omega j}$ inside the choice probabilities in Equation 8 restricts the choice set to schools where admission was feasible. The MLE therefore recovers the semi-elasticity of school choices to travel times (ϕ_g^S) subject to the restrictions imposed by the admission policy. Following Bayer et al. (2007), the estimation nests an inner fixed point that jointly recovers the transformed school quality $\tilde{S}_{jg} \equiv S_{jg}^{\theta_g^S}$. The semi-elasticity ϕ_g^S is identified by how strongly parents favor nearby schools. Differences by skill are captured if households in locations with different high-skill shares are more or less likely to choose closer schools. School quality is identified via revealed preferences, so schools that draw more students than the spatial distribution of households would predict must be of higher quality.

To separate the estimated semi-elasticity into its components, I again use a variance-matching procedure. I take standardized test scores of elementary-school students for 2012-2014, measured in the year before the application to reflect the information available to parents, average them at the school level, compute their variance within the municipality of Madrid, and match it to the variance of the school quality recovered in the estimation. I keep only schools with both pre-school and elementary stages, since in those cases parents choosing at the pre-school level likely already account for the quality of the elementary stage. The

¹⁸ Appendix B.2 gives further detail on the construction of the dataset.

¹⁹ Households with a child aged two or younger are those about to apply for the first year of pre-school, so their skill composition is the relevant one for the applicant pool. I measure $\hat{h}_{\omega}$ as the share of college-educated households among this demographic group in the applicant’s *sección*.

Table 2: School choice parameters

	ϕ_g^S	θ_g^S	κ^S
	(1)	(2)	(3)
	MLE	Var Match	Ratio
High-skill	0.432*** (0.003)	8.775	0.049
Low-skill	0.376*** (0.002)	7.642	0.049

Notes: Column (1) reports the MLE estimates for the sample of students admitted to their preferred option in the first year of pre-school for 2015; the text describes the remaining sample restrictions. Standard errors in parentheses. Columns (2) and (3) report the variance-matching estimates, which find the parameter that equates the variance of school quality to that of standardized test scores across primary schools that also offer a pre-school stage.
 $*p < 0.10$; $**p < 0.05$; $***p < 0.01$.

underlying assumption is that the dispersion of test scores across schools is representative of the dispersion of perceived quality, so that test results are strongly correlated with what parents value when choosing a school. As with commuting, I assume κ^S does not vary with skill, since the travel-time measure already does.²⁰ The procedure minimizes the distance between the variance of standardized tests and that of the recovered school-quality vectors:

$$\min_{\theta_H^S} \left\{ \text{Var}[\text{Test}_j] - \text{Var} \left[\frac{K_{jH}}{K_j} \tilde{S}_{jH}^{1/\theta_H^S} + \frac{K_{jL}}{K_j} \tilde{S}_{jL}^{1/\theta_L^S} \right] \right\}^2, \quad (22)$$

where the dispersion of standardized tests disciplines the dispersion of the school quality perceived by parents, as recovered from revealed preference in the MLE step.

Table 2 reports the estimates for 2015 and reveals that parents are highly sensitive to both school quality and travel times.²¹ The semi-elasticity of travel cost to travel time, κ^S , is 0.049, roughly 2.5 times the commuting value. The high value reflects the strong spatial concentration of school trips and points to a higher cost of traveling with children to school than of commuting alone to work. The dispersion parameters θ_g^S are also larger than their workplace counterparts, so quality weighs more heavily on school choice than wages do on labor supply, and idiosyncratic shocks matter correspondingly less. Finally, as an external check on the recovered quality, the quality measure correlates positively with schools' average test performance (Figure A.4). Only the second moment of test scores was used in the estimation, so the

²⁰The assumption implies $\theta_L^S = \theta_H^S \frac{\phi_L^S}{\phi_H^S}$, so that the dispersion of the single vector of standardized tests across schools identifies one common parameter from the two skill-specific school-quality vectors recovered in the MLE step.

²¹I focus on 2015 because it is the most recent year of data, and because the admission point system changed slightly from 2013 onward. The MLE estimates are stable across years, while the variance-matching estimates shift somewhat in 2014, tracking a different dispersion in test scores. Table A.2 reports the results by year.

untargeted correlation indicates that the revealed-preference measure captures a real dimension of quality, even if perceived quality also bundles many other school characteristics.

5.3 Residential choice, sorting and peer effects

Having established how households choose workplaces and schools conditional on their residence, I now estimate preferences for residential locations. Equation 12 shows that the choice depends on amenities, house prices, and workplace and school access. The key empirical challenge is that each element may be endogenous to households' residential choices.

As a first step, I compare households with and without children.²² As the commuting data show, workplace choices do not differ by parenthood, so workplace access is geographically equivalent for the two groups. Thus, taking the ratio of Equation 12 across the two groups removes workplace access. Next, I show that the comparison does not depend on house prices either. Table D.5 reports the effect of skills and parenthood on the housing expenditure share, using the Spanish Household Expenditure Survey.²³ Conditional on a few controls, two otherwise-equivalent households differ in their housing expenditure only by skill. Parenthood does not significantly affect the expenditure share, so the predicted α_{gk} collapse to α_g and house prices drop from the comparison:

$$\ln \frac{\lambda_{ng1}^B}{\lambda_{ng0}^B} = \ln \frac{\Phi_{g0}}{\Phi_{g1}} + \theta_g^B (\alpha_{g1} - \alpha_{g0}) \ln(r_n) + \theta_g^B \ln S_{ng} + \theta_g^B \ln \left(\frac{B_{ng1}}{B_{ng0}} \right) = \ln \frac{\Phi_{g0}}{\Phi_{g1}} + \theta_g^B \ln S_{ng} + \theta_g^B \ln \left(\frac{B_{ng1}}{B_{ng0}} \right). \quad (23)$$

The difference in residential choices between households with and without children is thus driven entirely by school access and amenities. I construct school access according to Equation 9 with parameters and quality recovered in the previous estimation step. The endogeneity of school quality in Equation 19, however, introduces a simultaneity problem: households sort into neighborhoods according to school access, the residential skill mix feeds into local school composition, which determines school quality, so access changes accordingly and drives further sorting. The econometrician cannot separate peer effects from residential preferences. For example, the concentration of high-skill households around high-quality schools may reflect either their stronger preferences or the response of school quality to their concentration.

I use two IV strategies to break the simultaneity between peer effects and residential sorting, one for each direction. The unobserved amenities and school quality fundamentals that drive the joint system are what the instruments must be excluded from. Each instrument severs one link of the causal chain, within the limits of the cross-sectional data.

Peer effects. I build a school-level dataset containing the recovered quality estimates, the skill composition of students, and other school characteristics to study peer effects. The aim is to quantify the effect of the

²²The approach is similar to Caetano (2019), who uses the location choices of households without children to absorb local amenities and isolate the value of school quality to parents. His identification assumes that parents and non-parents value amenities equally. Here, the comparison cancels only workplace access and house prices, and allows households to value unobserved amenities differently by parenthood, which requires quasi-experimental variation to recover residential preferences.

²³See Appendix D.5 for details on the estimation.

student mix on the quality parents perceive at the time of application, i.e. the elasticity δ_g^S in Equation 19. The mix is endogenous, however, as it is tied to residential sorting. Since the school-choice estimation shows that students come mostly from nearby locations, the mix correlates with fundamentals of school quality that do not depend on the mix itself and attract parents to live nearby, potentially differently by their skill level.

To recover an unbiased estimate of the peer-effect elasticity, I instrument the student mix with the school's elevation above sea level. The instrument's relevance follows from the history of Madrid's expansion (Pallol Trigueros, 2015). In the spirit of Heblich et al. (2021) and Loumeau (2023), who use historical pollution to explain sorting, the argument is that before the rollout of modern sewage systems, waste waters flowed from the higher parts of the city down toward the Manzanares, making the low-lying areas unpleasant and unhealthy and inducing sorting. As the city expanded outward from its historical core beginning in the 1870s, the northern and more elevated areas were settled by the affluent and the nobility, while the southern areas became slums and industrial districts. The sorting effects within the city persist today. Figure D.2 Panel A shows that the elevation still predicts the ratio of high-to-low-skill resident households. The exclusion restriction therefore assumes that elevation affects the skill mix only through the historical growth of the city and is unrelated to school fundamentals today. Panel B in Figure D.2 maps the elevation of each *sección*, showing that the elevation range is limited and that the gradient is likely inconsequential today due to modern transport and sewage systems.

Table 3 reports the peer-effects estimates.²⁴ Across all specifications, the coefficients for both high- and low-skill quality are positive and significant. In the preferred specification, Column (5), the elasticities are similar across groups, with $\delta_H^S = 0.038$ just above $\delta_L^S = 0.031$. Both are small, precise, and statistically close. A one-log-point rise in the high-to-low-skill ratio lifts perceived quality by three to four percent. Peer composition accounts for 8.9% of the residual variation in high-skill quality and 5.4% in low-skill quality, so most of what makes a school good lies elsewhere.

Identification rests on the exclusion restriction, that elevation shapes school quality only through the sorting it once produced. I control for the channels by which elevation might act directly, construction year, distance to the center, school type and price tier, size, and pedagogical offerings, and the estimate holds steady across them. What remains is an unobserved channel correlated with elevation given these controls, and two features make it unlikely. Madrid's public schools are funded uniformly, with per-pupil spending set regionally, so hidden funding should not correlate with elevation. And teacher quality, though unobserved, is unlikely to sort with elevation, since Madrid assigns teachers largely by seniority as posts open, leaving them little room to choose where they teach.

For high-skill quality the IV estimate is lower than OLS, a sign that sorting biases OLS upward. High-skill families choose good schools for features unrelated to peer composition, so the raw correlation overstates the true effect. The travel-survey evidence points to similar sorting. For low-skill quality the two estimates nearly coincide, since low-skill families sort more weakly. In both cases, because skill is approximated from

²⁴See Appendix D.6 for details on the data and sample of schools used, details about the instrumental variable and first stage.

Table 3: Peer effects in school quality

<i>Panel A: High-skill quality</i>						
	OLS			IV		
	(1)	(2)	(3)	(4)	(5)	(6)
Log skill ratio	0.054*** (0.005)	0.061*** (0.005)	0.055*** (0.005)	0.044*** (0.013)	0.038*** (0.012)	0.040*** (0.012)
Controls	No	Yes	Yes	No	Yes	Yes
Size			≥ 15			≥ 15
F-stat				81.03	100.84	97.41
Obs.	514	433	365	514	433	365
<i>Panel B: Low-skill quality</i>						
	OLS			IV		
	(1)	(2)	(3)	(4)	(5)	(6)
Log skill ratio	0.024*** (0.006)	0.030*** (0.006)	0.024*** (0.007)	0.034** (0.015)	0.031** (0.014)	0.034** (0.014)
Controls	No	Yes	Yes	No	Yes	Yes
Size			≥ 15			≥ 15
F-stat				81.03	100.84	97.41
Obs.	514	433	365	514	433	365

Notes: The unit of observation is a single school. The outcome is log school quality recovered from the school-choice estimation. The regressor is the log ratio of high- to low-skill admitted students. Columns (1)-(3) report OLS coefficients, Columns (4)-(6) instrument the skill ratio with elevation above sea level. All specifications weight schools by their total admitted students. Columns (3) and (6) also restrict to schools with at least 15 admitted students, to reduce measurement error in the ratio from small schools. Controls include construction year, distance to the center, school ownership, religious status, price tier, and pedagogical and language offerings. Heteroskedasticity-robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

the residential *sección*, measurement error in the student ratio pulls the bias the other way.

Residential choice. I estimate residential preferences by exploiting variation in school access across residential *secciones*, comparing households with and without children, as in Equation 23. From the Padrón Municipal, I compute the relative likelihood that households with and without children locate in each *sección*, the dependent variable. The main regressor is school access at the *sección* level, built from Equation 9 and the parameters and school quality recovered in previous steps.

Residential preferences face the same simultaneity as peer effects, running in the opposite direction. I therefore instrument school access with the average construction year of the schools around each *sección*. Importantly, I take the year of original construction rather than of any later renovation, which is more likely to respond to the same neighborhood dynamics the instrument is meant to avoid. School construction year predicts school quality strongly, plausibly because parents value newer buildings (Figure D.3).

Table 4: Residential preferences and school access

<i>Panel A: High-skill households</i>					
	OLS			IV	
	(1)	(2)	(3)	(4)	(5)
Log school access	-1.136*** (0.149)	0.199* (0.110)	0.171 (0.161)	7.822*** (1.842)	4.309** (1.766)
Controls		Yes	Yes	Yes	Yes
Barrio FE			Yes		Yes
F-stat				20.49	10.24
Obs.	2,245	2,244	2,240	2,222	2,218
<i>Panel B: Low-skill households</i>					
	OLS			IV	
	(1)	(2)	(3)	(4)	(5)
Log school access	-0.897*** (0.139)	-0.311*** (0.117)	-0.387*** (0.129)	5.191*** (1.118)	6.745** (2.927)
Controls		Yes	Yes	Yes	Yes
Barrio FE			Yes		Yes
F-stat				39.92	8.22
Obs.	2,245	2,244	2,240	2,222	2,218

Notes: The unit of observation is the residential *sección*. The outcome is the log ratio of the location shares of households with and without children, with the age cutoff of the youngest member set at 17 years old. The regressor is log school access, built from the school-choice estimates. Columns (4)-(5) instrument school access with the average construction year of schools within 1.5km weighted by inverse distance. Controls include the average dwelling size and building age within the *sección*, elevation and distance to the center. Columns (3) and (5) add fixed effects for the 128 *Barrios* of the municipality of Madrid. Heteroskedasticity-robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The exclusion restriction holds if the age of nearby schools affects residential choice only through school access, and not through amenities that households with and without children value differently. Because building age correlates with other neighborhood characteristics, I control for distance to the center, elevation above sea level, and the average construction year of all buildings in the *sección*, so the instrument exploits school-building timing net of the neighborhood's. I also control for average dwelling size, since households with and without children have different space needs that could sort them across locations by the housing stock. Under the exclusion restriction, the estimate does not capture the reverse channel by which high-amenity locations draw high-skill families and thereby raise school quality through peer effects. In the most demanding specification I add fixed effects for the 128 *barrios*, reducing the influence of unobserved amenities at the local level.

Table 4 presents the estimates.²⁵ Column (4), my preferred specification, instruments school access while conditioning on the observable confounds, and the first stage is strong. School access raises the relative presence of families for both groups, with $\theta_H^B = 7.8$ and $\theta_L^B = 5.2$. Column (5) adds *barrio* fixed effects, a more demanding test that absorbs unobserved local amenities. The fixed effects wash out much of the instrument’s variation; the first stage weakens materially for low-skill households and remains only modest for high-skill households, so confidence in the point estimates is lower in this specification. The coefficients stay positive, but given the weak first stage they offer only suggestive evidence against unobserved amenity confounding at the *barrio* level.

The OLS estimates are negative without controls and remain materially negative for low-skill households even after controls are added, although the magnitude falls relative to the raw specification. Two forces likely bias the OLS coefficients downward. As I discuss in Section 6, school access is highest in the center because of the density of schooling options. Here, families are less likely to reside, so school access is negatively related to the residual amenities that families value relative to non-families. School access is also a generated regressor, and the resulting measurement error attenuates OLS toward zero. The instrument removes both the confounding and the attenuation, driving the estimates above OLS.

5.4 Calibrated parameters

The remaining parameters are calibrated outside the estimation. I take the labor share $\beta = 0.8$ from Ahlfeldt et al. (2015), following Valentinyi and Herrendorf (2008), and the elasticity of substitution between skill groups $\rho = 1.43$ from Card (2009), in line with Tsivanidis (2026). The agglomeration elasticity is $\delta^A = 0.08$, as estimated for Berlin by Ahlfeldt et al. (2015). Housing supply elasticities come from Baum-Snow (2023), which imply supply parameters of $\mu_C = 0.265$ for the dense center, defined as the areas already built up by the 1960s, and $\mu_S = 0.371$ for the suburbs. Population stocks by skill and parenthood come from the Padrón for Madrid, averaged between 2013 and 2015. The city comprises 2,420 *secciones* and 514 schools. Table A.3 lists all the model parameters.

5.5 Model inversion

With the parameters in hand, I invert the model’s equilibrium conditions to recover the fundamentals that rationalize the observed data in Madrid. The recovery is recursive, and each fundamental is obtained in closed form from the observed allocations and the fundamentals recovered before it. Six values are obtained in the following order: the skill intensity ψ_i , the floorspace wedge ξ_n , the productivity fundamental v_i^A , land productivity ζ_n , the school-quality fundamental v_{jg}^S , and residential amenities B_{ngk} .

The model inversion concludes the model estimation. The preference and technology parameters are estimated from the choices of households and firms, the remaining parameters are calibrated from external sources, and the inversion recovers the location and school fundamentals that rationalize the observed city.

²⁵The appendix reports *sección*-level descriptive statistics (Table D.8), the first stage (Table D.9), and robustness to the family definition and to alternative instruments (Table D.10).

With the estimated model in hand, I next turn to what it reveals about the spatial structure of Madrid, and then to the counterfactual exercises.

6 Understanding school and workplace access through their margins

In this section, I introduce a new decomposition to the quantitative spatial literature, which allows separating any market access measure into two margins, an extensive and an intensive one. The separation is central to understanding what determines the different spatial distributions of school and workplace access. I then discuss what the decomposition implies for goods not allocated through prices, such as schools, and use it to anticipate the impact of introducing catchment areas.

Market access is the ex-ante expected utility from a discrete choice, and it depends on both how many options a household can reach and how good they are, its extensive and intensive margins.²⁶ The extensive margin counts the effective number of options through the inverse Herfindahl index of choice probabilities, which measures the number of equally attractive options that would generate the same concentration of choices (Adelman, 1969). A larger effective number of options is better, because it raises the chance that a household can choose an option it happens to like, as captured by a high idiosyncratic preference draw. The intensive margin, instead, captures the average utility of all options, weighted by the likelihood that they get chosen. Choosing among higher-value options is better, because it raises the expected utility of the one that gets selected. More generally, the decomposition rewrites access through choice probabilities, so both margins focus on the viable options while giving little weight to those that are rarely chosen. The separation into margins is exact and is derived in Appendix D.9.

To show how each margin empirically contributes to the spatial distribution of school and workplace access in Madrid, I report how they vary with distance to the center in Table 5.²⁷ I start with the value of options, in the first two columns of each panel. School quality rises slightly with distance, both for high- and low-skill households. However, the low R^2 signals that the gradient explains little of the variation in quality, and that at every distance from the center one finds schools of varying quality. Wages, on the other hand, tell a substantially different story, with a strong negative gradient that accounts for far more of the variation. The most central locations offer the highest wages, which fall as one moves away from the center.

Although school quality and wages move in opposite directions with distance to the city center, both access measures decline with it. In other words, while wages and workplace access are both highest in the center, school access peaks there even though school quality does not.²⁸ The decomposition of access explains why by revealing that the two measures follow separate margins, as shown in the rest of Table 5.

²⁶Throughout, I refer to the value of an option as what it offers before the trip, such as school quality or the wage, and to its utility as its value net of travel costs.

²⁷Figure A.5 in the Appendix plots the same relationships, and Figure A.6 includes maps of each access measure and its margins across *secciones* of the city.

²⁸The negative gradient of school access connects back to the estimation of residential preferences. Because it is on average higher in the center, where families are least likely to live, the OLS estimates turn out to be negative (see Table 4). In the model, other amenities specific to parents then explain the sorting observed in the data.

Table 5: Access gradients and their margins

<i>Panel A: Schools</i>								
	Quality		Access		Extensive		Intensive	
	High	Low	High	Low	High	Low	High	Low
Log dist. center	0.031** (0.013)	0.037*** (0.014)	-0.054*** (0.005)	-0.052*** (0.006)	-0.056*** (0.004)	-0.049*** (0.004)	0.002 (0.007)	-0.003 (0.008)
Obs.	514	514	2,413	2,413	2,413	2,413	2,413	2,413
R ²	0.011	0.015	0.098	0.061	0.099	0.076	0.000	0.000
<i>Panel B: Work</i>								
	Wages		Access		Extensive		Intensive	
	High	Low	High	Low	High	Low	High	Low
Log dist. center	-0.218*** (0.014)	-0.135*** (0.012)	-0.113*** (0.003)	-0.119*** (0.003)	0.039*** (0.001)	0.028*** (0.001)	-0.152*** (0.004)	-0.148*** (0.004)
Obs.	2,419	2,419	2,419	2,419	2,419	2,419	2,419	2,419
R ²	0.088	0.054	0.519	0.532	0.201	0.099	0.427	0.390

Notes: The unit of observation is a school for school quality, otherwise it is the *sección*. Each column regresses the log of the variable on log distance to the center (Plaza del Sol). The extensive and intensive margin elasticities sum up to the access ones. Heteroskedasticity-robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Workplace access declines outwards because of its intensive margin, which more than accounts for its gradient, while the effective number of workplaces actually rises with distance. The options reached from the center offer a much higher average utility, reflecting the steep gradient displayed by wages in the first place. School access, in turn, declines away from the city center only through its extensive margin, as the effective number of options drops toward the suburbs while the average school quality remains flat and statistically indistinguishable from zero.

Three key ingredients explain why the two measures follow different margins: travel costs, the elasticity of choices to utility, and the spatial distribution of value across options. First, the semi-elasticity of travel costs κ governs responsiveness to distance, and so how many options lie within reach and are actually considered. Second, the Fréchet parameter θ , i.e., the elasticity of choices to utility, sets how strongly differences in utility across options, whether from their value or travel costs, shift choice probabilities. Finally, the spatial distribution of the value of options determines the size and value of the effective choice set households obtain from their residential location.

In Madrid, the estimated ingredients of school and workplace access differ sharply. Both κ^S and θ_g^S exceed their workplace counterparts, so school choices fall quickly with travel time and school access remains very localized.²⁹ On top of that, as school quality is equally scattered throughout the city, the intensive margin shows no monocentric gradient, while the extensive margin displays a negative one, since schools

²⁹See Table A.3 for the comparison of the preferred estimates.

are denser in the center. By contrast, the much lower κ^W keeps many jobs across the city within reach, and even though the lower θ_g^W makes choices less responsive to wage differences, the steep gradient of wages drives workplace access through the intensive margin. Accordingly, the two measures also differ in how much of their variation is explained by distance to the city center. It accounts for about half of the variation in workplace access, against less than a tenth for school access. Workplace access is therefore smooth and centralized, while school access varies both locally across nearby neighborhoods and globally with distance to the center (see Figure A.6).

The decomposition also speaks directly to goods not allocated through prices, such as schools, where the allocation rule becomes a fourth ingredient of access. Admission probabilities scale the utility of each option, thus reshaping the concentration of choice probabilities and acting on both margins. Catchment areas are an allocation rule of exactly this kind, acting directly on school access. By setting the admission probability to zero for every school outside a household's catchment, they cut off the many options within reach that lie beyond its borders. The loss moves both margins, as it shrinks the effective number of options and shifts average utility depending on the quality of the lost schools.

In Madrid, the decomposition anticipates that catchment areas hit hardest in the center. There, the extensive margin is largest, and as catchments flatten the number of options across the city, removing the options outside them cuts access the most. The intensive margin adds a loss only where average quality differs systematically between catchment areas, so that borders separate households from better schools. Systematic quality differences matter little in Madrid, but they likely matter in many cities with uneven financing rules.

Unlike the intensive-margin one, the central extensive-margin effect should arise under most catchment policies, as it mostly depends on the geography of schooling options, and the conditions that sustain it are likely common to other cities beyond Madrid. First, the density of options must decline with distance to the center, as it does wherever population and child density fall away from it and school supply follows where children live. Second, willingness to travel to school must be low relative to how fast that density falls, which is especially likely in congested big cities where traveling is time consuming or where busing serves local schools only. Third, the variation in local quality must not differ much with distance to the center. Otherwise, once catchments bind, parents in the outskirts, for instance, could concentrate on a few good schools within their catchment while parents elsewhere continue to spread their choices, so that the effective number of options drops more where quality varies most. Finally, central options must not be much harder to get into, as they would be if central schools were systematically oversubscribed, which is unlikely where school supply adjusts to demand.³⁰ The case of workplace access in Madrid shows what happens when the conditions fail. Since commuting is cheap relative to how fast job density falls, and since the high-paying core makes the utility of jobs within reach far more unequal near the center, the workplace extensive margin actually rises with distance. By contrast, wherever the conditions hold for schools, the extensive margin is highest in the city center, and catchment areas that flatten choice sets cut

³⁰The school access measure derived in the paper already incorporates the admission probabilities under free choice, and its extensive margin still peaks in the center, so the condition holds for Madrid.

access most there, which lends external validity to the direction of the counterfactual results presented next.

7 Quantifying the effects of school catchment areas

The counterfactual I run substitutes free-choice school admission with a catchment-area policy, as is common in many cities around the world. To implement it, I restrict enrollment to the schools within residential *barrios*, one of Madrid’s own administrative units.³¹ A few *barrios* contain no school, so they are matched to the nearest one. The choice of geographical unit follows from the ease of using an existing administrative boundary, but it also corresponds to a size that effectively reduces the number of choices available to families.³²

I present the counterfactual outcomes by comparing the catchment-area policy against the free-choice baseline throughout, reporting every outcome in log changes relative to that baseline. Because the policy’s effect on school access depends on how strongly quality responds to the resulting sorting, I solve the same counterfactual at three values of the peer-effect elasticity δ_g^S : its estimated value, zero, and three times the estimate.³³ To isolate which margin of the city’s adjustment matters most, I further solve two variants that hold fixed, in turn, the supply of floorspace and its split between residential and commercial use, and the location choices of households without children.


The exercise carries three limitations worth stating. First, school capacity is not modeled, so admission is never rationed by available seats. The exercise thus requires a long-run interpretation, under which the local school administration adjusts capacity to demand. Consequently, the reported welfare changes understate the short-run cost of congestion at oversubscribed schools. Second, and for the same reason, the policy does not uniformly reduce options. In *secciones* where a *barrio*’s full roster of schools exceeds the effective number of options households enjoyed under free choice, catchment can increase the number of schools available (this happens for approximately 30% of *secciones*, as shown in Figure A.8). Third, the model is static and ignores any cost attached to relocation. To sum up, the results should be read as the policy’s long-run effect, describing how Madrid would look had it adopted catchment areas in the past.

³¹Effectively, I replace the estimated admission probability matrix between *secciones* and schools with one that has zeros for every out-of-zone school and ones for those inside. Madrid is divided into 128 *barrios*, and this division induces each *sección* to be linked to an average of 5.25 schools. Table A.4 reports the full distribution, using either *barrios* or *distritos* as catchment areas. See Figure A.7 for a map of the boundaries of *barrios* and *secciones*, and school locations.

³²Zone size poses a well-known trade-off for policymakers. Catchment areas are typically imposed to reduce parents’ travel and busing costs, manage school capacity, and ease enrollment planning. A zone that is too small removes choice almost entirely and ties school assignment directly to residential sorting. A zone that is too large, by contrast, achieves little where households already choose nearby schools, and so fails to constrain them at all. Figure A.8 shows the reduction in the number of choices under *barrio* catchment areas, compared with the larger *distrito*-based ones.

³³At the estimated δ_g^S , peer composition accounts for 8.9% of the variation in school quality for high-skill households and 5.43% for low-skill ones. Tripling the elasticity raises these shares to 80.1% and 48.9% respectively.

Figure 4: The effect of catchment areas on school access and residential choice gradients.



Notes: Each panel plots the log change in the variable against distance to the center, using a local polynomial smooth, comparing the equilibrium under *barrio* catchment against the free-choice baseline. Panel (a) reports the change in total school access and its two margins. Panel (b) reports the change in residential choices for each household type. Tables A.5 and A.6 report the corresponding linear gradients in regression form, also across scenarios.

7.1 Catchment areas and residential sorting

The effect of the catchment restriction on school access depends on the relevance of its two margins. Panel (a) of Figure 4 shows how catchment areas reduce the contribution of the extensive margin at every distance from the center. The drop is largest in the center, where the advantage of having more options under free choice accrued. The intensive margin, instead, rises everywhere, with the opposite gradient.³⁴ Its effect is nonetheless only marginal, so the overall change in school access, like the extensive margin driving it, is sharpest near the center. On average, a *sección* loses around 10% of its school access, markedly more so centrally.

Once school access is reshaped by the admission policy, parents respond directly by relocating, and other households indirectly react because of general equilibrium forces. Panel (b) of Figure 4 reports the residential relocation of households in response to the change in school access. Parents move the most, and follow closely the change in the gradient of school access, reallocating from the center to the peripheral areas of the city. The largest change is near the center, where the residential choice of high-skill parents falls by about 2.5 log points, roughly 92%. Non-parents adjust in the opposite direction, as a consequence of the changing prices in the city. Their response is only partial, at around 12% of the magnitude of parents' (see Table A.6), and so does not fully counteract the relocation of households with children. Nonetheless, accounting for their endogenous response turns out to be important in explaining the full consequences for the city, as I show in the next subsection.

Household relocation reshapes segregation both across neighborhoods and across schools. High-skill households self-select more strongly into areas with better access and then into better schools, and peer

³⁴The average *sección* gains in terms of the intensive margin because the options dropped outside of catchment are the same ones that had the lowest utility, i.e. school quality net of travel costs, to begin with.

effects further reinforce this mechanism.³⁵ The city-center gradients of Figure 4 mask these consequences, so I report the citywide figures in Panel B of Table 6, across different scenarios in its columns. The 2% citywide increase in residential segregation by skill is fully due to parents relocating to different peripheries, as segregation within this group grows by 8.6% while non-parents' own segregation does not change. School-level segregation moves far more, rising by 80.4% at the estimated peer-effect elasticity.³⁶ Both figures intensify sharply once peer effects are tripled, with residential segregation for parents rising to 30.3% and school segregation exploding to 214.2%. The comparison between results shows that when school quality depends more strongly on local demographics, through peer effects, catchment areas add a reinforcing channel to segregation.

The policy change moves households with and without children in opposite directions, increasing their segregation. It increases by 145.8%, more than ten times the increase in skill segregation. Again, this mechanism is reinforced by the intensity of peer effects, as school quality becomes more unequal across neighborhoods and incentivizes parents to relocate farther. The results however also underscore the importance of the endogenous relocation of non-parents. Roughly a third of the effect is driven by them, as holding this group fixed increases segregation by 98.7%. Overall, the city becomes more segregated along most dimensions, as parents move toward areas that guarantee better access, with high-skill households self-selecting into those offering the very best.

In terms of welfare, parents come out as clear losers from the policy, as they bear the cost of the school choice loss. Catchment areas reduce welfare by 7.4% for high-skill parents, and by 7.6% for low-skill ones, while leaving non-parents with a slim gain. Across the different scenarios presented in Panel A of Table 6, these figures remain remarkably stable, because the extensive margin, which mainly determines the school access change, is not endogenous to these channels (see Table A.5). The varying extent of relocation across scenarios also adds little welfare on top of it, consistent with households already being optimally sorted to begin with. The underlying mechanisms behind the catchment area effect are thus more relevant for the structure of the city, i.e. sorting and segregation, than for welfare.

7.2 The citywide effects of catchment areas

The structure of the city reacts to the introduction of catchment areas through general equilibrium forces in housing and labor markets. In practice, once the policy reshuffles residential demand, floorspace prices and wages adjust in turn. Figure 5 plots the gradient changes in both against distance to the center. Each panel reports three scenarios: the full model, one that keeps floorspace fixed at baseline, and one that does not allow non-parents to change residential locations. The comparison highlights the importance of these two adjustment mechanisms for the citywide effects.

Residential demand shifts outward as parents follow the change in school access. This first induces an

³⁵Segregation is measured through the exposure-gap index, as in [Diamond and Gaubert \(2022\)](#). It records the gap between the average local high-skill share experienced by high-skill households and by low-skill households, relative to the latter. I compute it separately for residential location and school enrollment, and also use it to compute residential segregation by parenthood.

³⁶Skill segregation within schools grows even when peer effects are shut down, signaling an existing correlation between the new distribution of school access and of neighborhood quality.

Table 6: The city-wide effects of catchment areas: aggregate results


	Full model	Fixed floorspace	Fixed non-parents	No peer effects	Peer effects $\times 3$
<i>Panel A: Welfare</i>					
High-skill parents	-7.39%	-7.80%	-7.60%	-7.95%	-7.11%
High-skill non-parents	0.23%	0.22%	0.45%	0.27%	0.09%
Low-skill parents	-7.63%	-7.95%	-7.73%	-7.58%	-9.29%
Low-skill non-parents	0.20%	0.09%	0.38%	0.15%	0.37%
<i>Panel B: Segregation</i>					
Households by skill	1.97%	1.46%	1.83%	-0.37%	9.51%
Parents by skill	8.64%	7.63%	8.32%	2.75%	30.26%
Non-parents by skill	-0.01%	0.25%	0.00%	0.30%	-1.07%
Households by parenthood	145.75%	141.75%	98.68%	133.02%	209.38%
Students by skill	80.40%	78.44%	79.92%	45.58%	214.24%
<i>Panel C: Labor market</i>					
Aggregate output	0.26%	-0.02%	0.43%	0.27%	0.20%
Aggregate productivity	0.07%	-0.03%	0.18%	0.08%	0.06%
Land-use dispersion	2.38%	0.00%	3.01%	2.54%	1.94%
High-skill parents wage	-0.05%	-0.27%	0.13%	-0.04%	-0.05%
High-skill non-parents wage	0.29%	0.09%	0.40%	0.34%	0.10%
Low-skill parents wage	0.08%	-0.29%	0.35%	0.07%	0.14%
Low-skill non-parents wage	0.41%	0.06%	0.61%	0.41%	0.44%
<i>Panel D: Travel times to work</i>					
High-skill parents	2.89%	2.66%	2.67%	3.08%	2.04%
High-skill non-parents	-0.46%	-0.54%	-0.05%	-0.45%	-0.46%
Low-skill parents	1.80%	1.58%	1.69%	1.78%	1.87%
Low-skill non-parents	-0.29%	-0.40%	0.06%	-0.29%	-0.29%
<i>Panel E: Travel times to school</i>					
High-skill parents	-43.46%	-43.87%	-43.23%	-43.54%	-43.13%
Low-skill parents	-39.49%	-39.93%	-39.23%	-39.55%	-39.24%

Notes: All entries are percentage changes relative to the free-choice baseline data. Fixed floorspace and fixed non-parents are counterfactuals that hold, respectively, the supply and split of floorspace and non-parents' location choices at their baseline values. No peer effects and peer effects $\times 3$ are counterfactuals that vary δ_g^S under the full model.

increase in floorspace prices, as supply cannot realign sufficiently. On top of that, as the aggregate population shifts outward from the city center, the gravity in commuting forces firms to incentivize workers to travel longer by increasing wages, thus reinforcing the baseline gradient steepness. Next, the long-dash line shows that the gradient change from fixing floorspace is 54% larger (see Panel B in Table A.7 for the linear coefficients), as fixing supply forces the market to clear through prices only. The labor market, however, barely responds under this scenario, as with floorspace pinned, the change in labor supply is rebalanced one-to-one by wage changes and the distribution remains the same.

The endogenous response of non-parents dampens the direct effect of the policy. The short-dash line

Figure 5: The effect of catchment areas on the price gradients.



Notes: Each panel plots the log change in the variable against distance to the center, using a local polynomial smooth, comparing the equilibrium under *barrio* catchment against the free-choice baseline. It plots three lines across scenarios: solid for the full model, long dash for fixed floorspace, and short dash for fixed non-parents. Panel (a) reports the gradient change in floorspace price; Panel (b) the gradient change in high-skill wages. The response of low-skill wages is very similar. Table A.7 reports the corresponding linear gradients in regression form.

traces an equivalently stronger reaction of floorspace prices, driven this time by the loss of non-parents' demand-side buffer, rather than by fixed supply. This result, however, should be interpreted with caution. As the model abstracts from endogenous amenities, the dampening role played by non-parents here is likely overstated. It is reasonable to imagine that the relocation of parents and skilled households might induce neighborhood change and draw non-parents to follow. The comparison similarly finds dampened effects for the labor market, as the relocation of non-parents towards the center allows firms to reduce the wage adjustment, via the substitution of workers from households with children to those without.³⁷ Overall, these results confirm that non-parents play an important role, and that reduced-form results that ignore them risk understating the impact of school policy on the target population of parents.³⁸


Spatially, catchment areas drive suburbanization and a stronger separation of residential neighborhoods from workplaces. The first panel of Figure 6 shows that total floorspace adjusts to the policy by densifying the suburbs. On top of that, the split between commercial and residential floorspace diverges following households' and firms' demand, as shown by the gradient changes in Panel C of Table A.7, and by the 2.4% increase in land-use dispersion across the two types.³⁹ Consistently, the remaining panels show the

³⁷The underlying assumption here is that the two groups are perfect substitutes, although from a life-cycle perspective workers with children are more likely to have higher experience and thus be more productive. So again, the dampening effect of non-parents needs to be interpreted with caution, although it is difficult to argue that it would be null.

³⁸For example, empirical investigations of the valuation of school quality that use boundary discontinuities around catchment areas might understate the valuation, because the house price effects would result from the GE response of non-parents, who may relocate to the lower-quality side to avoid the price increases.

³⁹Land-use dispersion applies the same exposure-gap index, with the share of residential floorspace in each neighborhood replacing the high-skill share. It records the gap between the average local residential share experienced by residential floorspace and by commercial floorspace, relative to the latter.

Figure 6: The effect of catchment areas on the gradients of total floorspace, residents, and workers.



Notes: Each panel plots the log change in the variable against distance to the center, using a local polynomial smooth, comparing the equilibrium under *barrio* catchment against the free-choice baseline. It plots three lines across scenarios: solid for the full model, long dash for fixed floorspace, and short dash for fixed non-parents. Panel (a) reports the gradient change in total floorspace; Panel (b) the gradient change in total households; Panel (c) the gradient change in the total workforce. Table A.7 reports the corresponding linear gradients in regression form.

residential population moving outward, while the workforce drifts in the opposite direction but less, as commuting lets workers relocate without changing workplace. Taken together, these patterns confirm that, by capping the extensive margin advantage of central locations, catchment areas noticeably reshape the city structure.

The opposite reallocation of residents and workers makes the city more productive, but requires parents to commute more. Central locations that were already more productive to begin with are now attracting more workers and enjoying stronger agglomeration.⁴⁰ The accompanying wage increases accrue across most worker types, but not for high-skill parents, who suburbanize the most (see Panel A of Table A.7). Accordingly, Panel D of Table 6 shows that the average commute for high-skilled parents rises the most, while non-parents now commute less. These results reinforce the commuting penalty for parents documented in the motivating facts (see Figure 3), while average travel times to school are cut by a large 40%, even though they were shorter than commutes to begin with.

Overall, weighing free-choice and catchment-area admission systems poses a trade-off for policy makers. The latter starkly reduce travel times to school, providing the most direct rationale for the policy, and the simplest metric to justify its success. Ignoring the reduction in school choice in the denser city center, however, hides the overall effects on parents and the city structure. First, the model shows how costs mostly fall on the families the policy claims to serve. Parents' welfare shrinks, as they lose schooling options, relocate farther from the city center, pay the higher house prices now prevailing there, and commute longer to work. Second, neighborhoods become markedly more segregated, by skill and parenthood, and

⁴⁰These gains depend on floorspace actually reallocating toward these locations, as shown in Panel C of Table 6. Absent that margin, the productivity gains reverse as workers slightly move away from the city center.

the same happens to schools. These distributional costs are large and grow significantly with the intensity of peer effects, so policy makers should carefully consider the school quality feedback. Against these results, the same reorganization of the city brings some gains. As locations specialize between residential and commercial uses, output and wages rise and, at the same time, non-parents commute less and their welfare improves slightly. Catchment areas, in this sense, induce incidental gains that could blunt political opposition, even though they seem modest relative to their distributional costs. Though the model is silent on this, support for catchment areas could be strong among parents, and especially high-skilled ones, if they have genuine preferences for segregation both at the neighborhood and school level.

8 Concluding remarks

This paper studies school access and evaluates how catchment areas for school enrollment reshape residential choices and labor market outcomes across the city. Introducing catchment areas ties school admission to residential location, so parents may need to trade off access to their children's school against access to their jobs, and their choices could spill over to households without children. To capture the full extent of the policy's effects, I develop a quantitative spatial model that links school choice to the housing and labor markets. Households choose where to live and where to work, and parents also choose which school their children attend. School quality responds to the endogenous student skill mix, and floorspace supply adapts to the demand for residential and commercial uses. I estimate the model for Madrid with cross-sectional data on school admissions, commuting flows, and residential choices. Exploiting two historical instruments in the estimation allows me to identify peer effects in school quality separately from residential sorting and preferences for neighborhoods.

I run a counterfactual that imposes catchment areas across the city, along the borders of existing administrative units. The key insight of the analysis comes from decomposing workplace and school access into their intensive and extensive margins, which explains why their spatial distributions differ and clarifies the consequences of catchment areas. I show that the extensive margin, the number of viable options to choose from, rather than their average quality, is the main driver of the variation in school access across neighborhoods. For this reason, catchment areas penalize central city locations by restricting choice where the density of options is highest, thereby cutting school access by about 10% on average. As a result, parents relocate away from the city center, and the city structure changes in response. In particular, residential segregation increases, as high-skill parents sort more strongly to access the best schools and non-parents move out of those same areas. Moreover, the adjustment of floorspace use allows neighborhoods to specialize as residential demand shifts outward, so that workers concentrate more in the city center, where productivity is highest, and overall output slightly grows.

Overall, this paper shows that evaluating catchment areas requires tracing their effects from schools to housing and labor markets, which the local effects documented in the literature typically ignore. Looking only at school trips, the policy succeeds, with travel times to school falling by 40%. However, the longer school trips under free choice reflect gains that parents obtain elsewhere, and the costs of restricting choice appear once the full stream of citywide effects unfolds. Parents end up with longer commutes to work

and bear sizable welfare losses of around 7.5%. At the same time, neighborhoods and schools grow more segregated, while the gains in productivity and non-parents' welfare remain modest. The results thus call for a reassessment of why so many cities adopt catchment areas for school choice.

References

- ABDULKADIOĞLU, A., N. AGARWAL, AND P. A. PATHAK (2017): “The Welfare Effects of Coordinated Assignment: Evidence from the New York City High School Match,” *American Economic Review*, 107, 3635–3689. (Cited on page 4.)
- ADELMAN, M. A. (1969): “Comment on the “H” Concentration Measure as a Numbers-Equivalent,” *The Review of Economics and Statistics*, 51, 99. (Cited on pages 2, 5, 26, and 67.)
- AGARWAL, N. AND P. SOMAINI (2018): “Demand Analysis Using Strategic Reports: An Application to a School Choice Mechanism,” *Econometrica*, 86, 391–444. (Cited on page 4.)
- AGOSTINELLI, F., M. LUFLADE, AND P. MARTELLINI (2024): “On the Spatial Determinants of Educational Access,” NBER Working Paper 32246. (Cited on page 4.)
- AHLFELDT, G. M., S. J. REDDING, D. M. STURM, AND N. WOLF (2015): “The Economics of Density: Evidence From the Berlin Wall,” *Econometrica*, 83, 2127–2189. (Cited on pages 1, 2, 5, 12, and 25.)
- ALBOUY, D. AND R. J. FABERMAN (2025): “Skills, Migration and Urban Amenities over the Life Cycle,” NBER Working Paper 33552. (Cited on page 5.)
- ALMAGRO, M. AND T. DOMÍNGUEZ-IÑO (2019): “Location Sorting and Endogenous Amenities: Evidence from Amsterdam,” Working paper. (Cited on page 5.)
- BAUM-SNOW, N. (2023): “Constraints on City and Neighborhood Growth: The Central Role of Housing Supply,” *Journal of Economic Perspectives*, 37, 53–74. (Cited on page 25.)
- BAYER, P., F. FERREIRA, AND R. McMILLAN (2007): “A Unified Framework for Measuring Preferences for Schools and Neighborhoods,” *Journal of Political Economy*, 115, 588–638. (Cited on pages 1, 4, and 19.)
- BENABOU, R. (1993): “Workings of a City: Location, Education, and Production,” *The Quarterly Journal of Economics*, 108, 619–652. (Cited on page 4.)
- BIBLER, A. AND S. B. BILLINGS (2020): “Win or Lose: Residential Sorting After a School Choice Lottery,” *The Review of Economics and Statistics*, 102, 457–472. (Cited on page 4.)
- BILLINGS, S. B., D. J. DEMING, AND J. ROCKOFF (2014): “School Segregation, Educational Attainment, and Crime: Evidence from the End of Busing in Charlotte-Mecklenburg,” *The Quarterly Journal of Economics*, 129, 435–476. (Cited on page 4.)
- BLACK, S. E. (1999): “Do Better Schools Matter? Parental Valuation of Elementary Education,” *The Quarterly Journal of Economics*, 114, 577–599. (Cited on pages 1 and 4.)
- BLACK, S. E. AND S. MACHIN (2011): “Housing Valuations of School Performance,” in *Handbook of the Economics of Education*, ed. by E. A. Hanushek, S. Machin, and L. Woessmann, Elsevier B.V., vol. 3, 485–519. (Cited on page 4.)

- BOGART, W. T. AND B. A. CROMWELL (2000): “How Much Is a Neighborhood School Worth?” *Journal of Urban Economics*, 47, 280–305. (Cited on page 4.)
- CAETANO, G. (2019): “Neighborhood sorting and the value of public school quality,” *Journal of Urban Economics*, 114, 103193. (Cited on page 21.)
- CAETANO, G. AND H. MACARTNEY (2021): “What determines school segregation? The crucial role of neighborhood factors,” *Journal of Public Economics*, 194, 104335. (Cited on pages 1 and 4.)
- CALDWELL, S. AND O. DANIELI (2024): “Outside Options in the Labour Market,” *Review of Economic Studies*, 91, 3286–3315. (Cited on page 5.)
- CALSAMIGLIA, C., C. FU, AND M. GÜELL (2020): “Structural Estimation of a Model of School Choices: The Boston Mechanism versus Its Alternatives,” *Journal of Political Economy*, 128, 642–680. (Cited on page 14.)
- CAMPOS, C. AND C. KEARNS (2024): “The Impact of Public School Choice: Evidence from Los Angeles’s Zones of Choice,” *The Quarterly Journal of Economics*, 139, 1051–1093. (Cited on pages 4 and 6.)
- CARD, D. (2009): “Immigration and inequality,” *American Economic Review: Papers and Proceedings*, 99, 1–21. (Cited on page 25.)
- COUTURE, V., C. GAUBERT, J. HANDBURY, AND E. HURST (2023): “Income Growth and the Distributional Effects of Urban Spatial Sorting,” *Review of Economic Studies*. (Cited on page 5.)
- DIAMOND, R. (2016): “The Determinants and Welfare Implications of US Workers’ Diverging Location Choices by Skill: 1980–2000,” *American Economic Review*, 106, 479–524. (Cited on page 5.)
- DIAMOND, R. AND C. GAUBERT (2022): “Spatial Sorting and Inequality,” *Annual Review of Economics*, 14, 795–819. (Cited on page 31.)
- ECKERT, F. AND T. KLEINEBERG (2021): “Saving the American Dream? Education Policies in Spatial General Equilibrium,” Working paper. (Cited on page 5.)
- EPPEL, D. AND H. SIEG (1999): “Estimating Equilibrium Models of Local Jurisdictions,” *Journal of Political Economy*, 107, 645–681. (Cited on page 4.)
- FEENSTRA, R. C. (1994): “New Product Varieties and the Measurement of International Prices,” *American Economic Review*, 84, 157–177. (Cited on pages 5 and 67.)
- FERNANDEZ, R. AND R. ROGERSON (1998): “Income Distribution and Public Education: A Dynamic Quantitative Evaluation of School Finance Reform,” *American Economic Review*, 88, 813–833. (Cited on page 4.)
- FOGLI, A., V. GUERRIERI, M. PONDER, AND M. PRATO (2026): “The End of the American Dream? Inequality and Segregation in US Cities,” *Journal of Political Economy*, 134, 1110–1158. (Cited on page 5.)
- GIBBONS, S., S. MACHIN, AND O. SILVA (2013): “Valuing school quality using boundary discontinuities,” *Journal of Urban Economics*, 75, 15–28. (Cited on page 4.)

- GORTÁZAR, L., D. MAYOR, AND J. MONTALBÁN (2023): “Residence-based priorities and school choice,” *Economics of Education Review*, 95, 102384. (Cited on pages 4 and 42.)
- GREAVES, E. AND H. TURON (2024): “School Choice and Neighborhood Sorting: Equilibrium Consequences of Geographic School Admissions,” Working paper. (Cited on page 4.)
- HAHM, D. W. AND M. PARK (2022): “Location Choice, Commuting, and School Choice,” Working paper. (Cited on page 5.)
- HANDBURY, J. AND D. E. WEINSTEIN (2015): “Goods Prices and Availability in Cities,” *Review of Economic Studies*, 82, 258–296. (Cited on pages 5 and 67.)
- HEBLICH, S., S. J. REDDING, AND D. M. STURM (2020): “The Making of the Modern Metropolis: Evidence from London,” *Quarterly Journal of Economics*, 135, 2059–2133. (Cited on page 5.)
- HEBLICH, S., A. TREW, AND Y. ZYLBERBERG (2021): “East-Side Story: Historical Pollution and Persistent Neighborhood Sorting,” *Journal of Political Economy*, 129, 1508–1552. (Cited on pages 2 and 22.)
- KAPOR, A. J., C. A. NEILSON, AND S. D. ZIMMERMAN (2020): “Heterogeneous Beliefs and School Choice Mechanisms,” *American Economic Review*, 110, 1274–1315. (Cited on page 4.)
- KUMINOFF, N. V., V. K. SMITH, AND C. TIMMINS (2013): “The New Economics of Equilibrium Sorting and Policy Evaluation Using Housing Markets,” *Journal of Economic Literature*, 51, 1007–1062. (Cited on page 4.)
- LAVERDE, M. (2020): “Unequal Assignments to Public Schools and the Limits of School Choice,” Working paper. (Cited on page 4.)
- LAVY, V. (2010): “Effects of Free Choice Among Public Schools,” *Review of Economic Studies*, 77, 1164–1191. (Cited on page 4.)
- LEE, K. H. AND B. J. TAN (2026): “Urban Transit Infrastructure and Inequality,” *Review of Economics and Statistics*, 1–18. (Cited on page 18.)
- LEE, Y. S. (2015): “School Districting and the Origins of Residential Land Price Inequality,” *Journal of Housing Economics*, 28, 1–17. (Cited on page 4.)
- LOUMEAU, G. (2023): “Locating Public Facilities: Theory and Micro Evidence from Paris,” *Journal of Urban Economics*, 135, 103544. (Cited on pages 2, 5, and 22.)
- MACHIN, S. AND K. G. SALVANES (2016): “Valuing School Quality via a School Choice Reform,” *Scandinavian Journal of Economics*, 118, 3–24. (Cited on page 4.)
- MIYAUCHI, Y., K. NAKAJIMA, AND S. J. REDDING (2025): “The Economics of Spatial Mobility: Theory and Evidence Using Smartphone Data,” *The Quarterly Journal of Economics*, 140, 2507–2570. (Cited on pages 2, 5, and 12.)

- MORENO-MALDONADO, A. AND C. SANTAMARIA (2021): “Delayed Childbearing and Urban Revival,” Working paper. (Cited on page 5.)
- NCEE (2026): “School Choice: A Global Snapshot,” NCEE monograph, National Center on Education and the Economy, Washington, DC. (Cited on page 1.)
- PALLOL TRIGUEROS, R. (2015): *El Ensanche Norte. Chamberí, 1860-1931. Un Madrid moderno*, Catarata. (Cited on page 22.)
- REBACK, R. (2005): “House prices and the provision of local public services: capitalization under school choice programs,” *Journal of Urban Economics*, 57, 275–301. (Cited on page 4.)
- RIES, J. AND T. SOMERVILLE (2010): “School Quality and Residential Property Values: Evidence from Vancouver Rezoning,” *Review of Economics and Statistics*, 92, 928–944. (Cited on page 4.)
- SANTOS SILVA, J. M. C. AND S. TENREYRO (2006): “The Log of Gravity,” *Review of Economics and Statistics*, 88, 641–658. (Cited on page 17.)
- SIEG, H., V. K. SMITH, H. S. BANZHAF, AND R. WALSH (2004): “Estimating the General Equilibrium Benefits of Large Changes in Spatially Delineated Public Goods,” *International Economic Review*, 45, 1047–1077. (Cited on page 4.)
- TIEBOUT, C. M. (1956): “A Pure Theory of Local Expenditures,” *Journal of Political Economy*, 64, 416–424. (Cited on page 4.)
- TSIVANIDIS, N. (2026): “Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogotá’s TransMilenio,” *American Economic Review*, 116, 418–463. (Cited on pages 2, 5, 12, 14, 18, and 25.)
- VALENTINYI, A. AND B. HERRENDORF (2008): “Measuring factor income shares at the sectoral level,” *Review of Economic Dynamics*, 11, 820–835. (Cited on page 25.)

Appendix: Table of Contents

A	Additional tables and figures	42
B	Data construction	53
B.1	Travel times	53
B.2	School admission data	54
C	Model derivation details	56
C.1	Firms	56
C.2	Floorspace supply	56
D	Additional estimation results	57
D.1	Workplace choice estimation	57
D.2	Admission probabilities and skills	58
D.3	Approximation of school access	58
D.4	School choice estimation	59
D.5	Estimation of expenditure share by skill and parenthood	60
D.6	Peer-effects estimation	61
D.7	Residential choice estimation	63
D.8	Inversion and recovery of fundamentals	66
D.9	Access decomposition into extensive and intensive margins	67


Appendix A: Additional tables and figures

Figure A.1: Table on the school admission point system from [Gortázar et al. \(2023\)](#).

BONUS	CRITERIA	NUMBER OF POINTS		
		Before 2012/2013	2012/2013	2013/2014
Proximity -Madrid city-	Family house or parents' work in:			
	School district	4	4	
	Boundary school district	2	2	
	Family house or parents' work in:			
	Same municipality			4
	School district			0.5
	Region of Madrid			2
Low-income	Income $\leq$ IPREM	2		
	IPREM <Income $\leq$ 2IPREM	1		
	Minimum Insertion Subsidy		2	2
Siblings	First one 4pts, and additional 3pts	4		
	One or more		8	10
Disability	Parents, students or siblings	1.5	1.5	1.5
Large Family	General	1.5	1.5	1.5
	Special	2.5	2.5	2.5
Alumni family member	Family member is alumni student		1.5	1.5
School discretionary		1	1	1


Notes: The changes beyond the proximity criteria were applied together across all medium and large municipalities. IPREM is the acronym in Spanish for the Multiple Effects Income Public Index, which was €7,455.14 in the period of study. The Minimum Insertion Subsidy (*Renta Mínima de Inserción*) is a special provision granted for people with lower income than IPREM. School discretionary is a point that the schools have freedom to assign based on “public and objective” criteria.

Figure A.2: Share of school admissions by grade.




Notes: The Figure reports the share of admitted students via the centralized system in the municipality of Madrid for years 2012-2015 across all grades. It highlights that most admissions happen at the point of entry, in the first year of preschool, as admitted students always have a guaranteed seat for the next year for compulsory schooling years.

Figure A.3: Trips by motive




Notes: Panel a: The graph reports the separate distributions of travel times for trips to a workplace or a school, as self-reported by respondents of the EDM survey. The dashed lines are set at the median value of each distribution. The distributions show the extent of the bias in reporting, with clear anchor points at 30 and 60 minutes, as well as other multiples of 5. Panel b: The graph reports the separate distributions of distances for trips to a workplace or a school. The dashed lines are set at the median value of each distribution. Distances are computed from Google Maps using car as travel mode, thus accounting for the road network, between origin and destination *barrios*.

Figure A.4: Recovered school quality and test scores




Notes: each panel plots the school quality recovered from the school-choice estimation against the school's average elementary standardized test score, for high-skill (Panel a) and low-skill (Panel b) households. Only the second moment of test scores is used in the estimation. The fitted line shows the correlation.

Figure A.5: Access and its margins with distance to the center



Notes: Each panel plots the log of the variable against distance to the center, with the scatter across *secciones* and a linear fit. Panel a reports the two underlying destination-level utility measures, school quality and wages, both in deviation from their mean. Panels b and c report each access measure together with its extensive and intensive margins, whose logs sum to log access in every *sección*. All variables are averaged across skill groups.

Figure A.6: Access and its margins across *secciones*




Notes: Each panel maps the log of the variable across the *secciones* of the municipality of Madrid, in deviation from its mean, with darker colors indicating higher values. The top row reports school access (Panel a) and its extensive (Panel b) and intensive (Panel c) margins; the bottom row reports the same decomposition for workplace access (Panels d to f). The extensive and intensive margins sum to access in every *sección*. All variables are averaged across skill groups. Gray areas are *secciones* with missing data.

Figure A.7: *Barrios* and *secciones* boundaries with school locations.



Notes: The map draws the boundaries of the 128 *barrios* of the municipality of Madrid, with their constituent *secciones* outlined in a thinner stroke. School locations are marked in green.

Figure A.8: Change in available options between free school choice and catchment areas.



Notes: This figure shows the change in the available schooling options between free choice, as computed through the effective number of options, and catchment areas. Negative values indicate that the zone restricts choices, positive ones that the zone doesn't bind. The *barrio* catchment area is on the left graph, the *distrito* one on the right. Skills depend purely on the effective number of options value for the two groups under free choice.

Table A.1: Spatial behavior by skill and parenthood

<i>Panel A: School trips by skills</i>				
	Travel time to school			
	(1)	(2)	(3)	(4)
High-skill HH	0.102*** (0.015)	0.049*** (0.017)	0.010 (0.016)	-0.034** (0.015)
Controls		✓	✓	✓
Home FE			✓	
Dest. FE				✓
Observations	6,179	5,536	5,536	5,534
<i>Panel B: Work trips by skill</i>				
	Travel time to work			
	(1)	(2)	(3)	(4)
High-skill HH	0.052*** (0.011)	0.017 (0.019)	0.020 (0.018)	0.000 (0.018)
Controls		✓	✓	✓
Home FE			✓	
Dest. FE				✓
Observations	17,275	12,590	12,590	12,590
<i>Panel C: Work trips by parenthood</i>				
	Travel time to work			
	(1)	(2)	(3)	(4)
HH with children	0.064*** (0.011)	0.025 (0.016)	0.013 (0.015)	0.031** (0.014)
Controls		✓	✓	✓
Home FE			✓	
Dest. FE				✓
Observations	17,275	12,590	12,590	12,590

Notes: The unit of observation is the individual trip from the EDM survey. Outcome is the log travel times, measured via Google Maps between origin and destination *Barrios*, averaging car, transit, and walking times with predicted weights by travel motive and skill. Column (1) has no controls; (2) adds individual and household controls; (3) adds residence (home) fixed effects; (4) adds destination fixed effects (school in Panel A, workplace in Panels B–C). Controls include age group, sex, education level, Spanish nationality, reduced mobility, household size, number of vehicles, carers, and workers, mode of transport, and, for workers only, job activity, sector, and occupation. Observations weighted by survey weights. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.2: School choice parameters by year

Year	ϕ_H^S	ϕ_L^S	θ_H^S	θ_L^S	κ^S
2012	0.441***	0.409***	–	–	–
2013	0.437***	0.382***	8.354	7.314	0.052
2014	0.431***	0.378***	10.410	9.133	0.041
2015	0.432***	0.376***	8.775	7.642	0.049

Notes: ϕ_g^S estimated by MLE for each year. θ_g^S and κ^S obtained by variance matching, which requires the prior year's test-score variance as the target moment; the 2012 row therefore reports ϕ_g^S only. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table A.3: Model parameters

Parameter	High skill	Low skill	Source
<i>Panel A: Workplace choice</i>			
θ_g^W	4.366	5.087	Commuting gravity regression
κ^W	0.019		
<i>Panel B: School choice</i>			
θ_g^S	8.775	7.642	School choice MLE
κ^S	0.049		
<i>Panel C: Peer effects</i>			
δ_g^S	0.038	0.031	Peer effects IV regression
<i>Panel D: Residential choice</i>			
θ_g^B	7.822	5.191	Residential choice IV regression
<i>Panel E: Expenditure shares</i>			
α_{g0}	0.372	0.438	Expenditure regression
α_{g1}	0.362	0.434	
<i>Panel F: Calibrated parameters</i>			
β	0.800		Valentinyi & Herrendorf, 2008
ρ	1.430		Card, 2009
δ^A	0.080		Ahlfeldt et al., 2015
μ_C	0.265		Baum-Snow, 2023
μ_S	0.371		Baum-Snow, 2023

Table A.4: Distribution of the number of available schools per *sección*, by catchment unit

<i>Barrio</i>			<i>Distrito</i>		
Schools	Freq.	Percent	Schools	Freq.	Percent
1	220	9.09	8	30	1.24
2	194	8.02	13	46	1.90
3	402	16.61	14	118	4.88
4	229	9.46	16	89	3.68
5	426	17.60	17	390	16.12
6	214	8.84	20	235	9.71
7	259	10.70	22	84	3.47
8	202	8.35	24	119	4.92
10	99	4.09	25	99	4.09
11	33	1.36	26	110	4.55
12	96	3.97	27	94	3.88
14	46	1.90	32	103	4.26
			36	171	7.07
			37	176	7.27
			38	178	7.36
			42	201	8.31
			46	177	7.31

Notes: The unit of observation is the *sección* (N=2420). Schools counts a *sección*'s admissible schools under each catchment unit, including the nearest-school fallback for zones without a school.

Table A.5: The city-wide effects of catchment areas: school access and its margins

<i>Panel A: $\Delta \ln$ School access</i>										
	Full model		Fixed floorspace		Fixed non-parents		No peer effects		Peer effects $\times 3$	
	High	Low	High	Low	High	Low	High	Low	High	Low
Log dist. center	0.047*** (0.006)	0.053*** (0.007)	0.047*** (0.006)	0.053*** (0.007)	0.047*** (0.006)	0.053*** (0.007)	0.047*** (0.006)	0.053*** (0.007)	0.049*** (0.007)	0.054*** (0.008)
Mean	-0.099	-0.098	-0.098	-0.098	-0.099	-0.098	-0.100	-0.099	-0.116	-0.113
Obs.	2,419	2,419	2,419	2,419	2,419	2,419	2,419	2,419	2,419	2,419
R ²	0.037	0.038	0.037	0.038	0.037	0.038	0.037	0.038	0.031	0.033
<i>Panel B: $\Delta \ln$ Extensive margin</i>										
	Full model		Fixed floorspace		Fixed non-parents		No peer effects		Peer effects $\times 3$	
	High	Low	High	Low	High	Low	High	Low	High	Low
Log dist. center	0.073*** (0.004)	0.070*** (0.004)	0.073*** (0.004)	0.070*** (0.004)	0.073*** (0.004)	0.070*** (0.004)	0.073*** (0.004)	0.070*** (0.004)	0.072*** (0.004)	0.070*** (0.004)
Mean	-0.137	-0.130	-0.137	-0.130	-0.137	-0.130	-0.137	-0.131	-0.138	-0.130
Obs.	2,419	2,419	2,419	2,419	2,419	2,419	2,419	2,419	2,419	2,419
R ²	0.143	0.125	0.143	0.125	0.143	0.125	0.145	0.125	0.139	0.125
<i>Panel C: $\Delta \ln$ Intensive margin</i>										
	Full model		Fixed floorspace		Fixed non-parents		No peer effects		Peer effects $\times 3$	
	High	Low	High	Low	High	Low	High	Low	High	Low
Log dist. center	-0.026*** (0.004)	-0.017*** (0.005)	-0.026*** (0.004)	-0.017*** (0.005)	-0.026*** (0.004)	-0.017*** (0.005)	-0.027*** (0.005)	-0.018*** (0.005)	-0.023*** (0.005)	-0.016*** (0.005)
Mean	0.045	0.037	0.046	0.038	0.046	0.038	0.045	0.037	0.025	0.019
Obs.	2,419	2,419	2,419	2,419	2,419	2,419	2,419	2,419	2,419	2,419
R ²	0.024	0.009	0.025	0.009	0.025	0.009	0.026	0.009	0.013	0.006

Notes: The unit of observation is the *sección*. Each column regresses the log change in school access, or one of its two margins, on log distance to the center (Plaza del Sol), against the free-choice baseline in the data. Full model is the estimated barrio catchment policy; fixed floorspace and fixed non-parents hold, respectively, the residential-commercial floorspace split and non-parents' location choices at their baseline values; no peer effects and peer effects $\times 3$ vary δ_g^S under the full model. The extensive- and intensive-margin coefficients in Panels B and C sum up to the access coefficient in Panel A, since the two margins multiply back to school access exactly (see Appendix D.9). Heteroskedasticity-robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.6: The city-wide effects of catchment areas: households and workers by type

<i>Panel A: Residential location</i>										
	Full model				Fixed floorspace				Fixed non-parents	
	Parents		Non-parents		Parents		Non-parents		Parents	
	High	Low	High	Low	High	Low	High	Low	High	Low
Log dist. center	0.329*** (0.049)	0.244*** (0.036)	-0.040*** (0.003)	-0.031*** (0.003)	0.310*** (0.047)	0.229*** (0.035)	-0.059*** (0.006)	-0.045*** (0.005)	0.308*** (0.048)	0.226*** (0.036)
Obs.	2,419	2,419	2,419	2,419	2,419	2,419	2,419	2,419	2,419	2,419
R ²	0.034	0.034	0.046	0.045	0.033	0.033	0.045	0.044	0.032	0.032
<i>Panel B: Workplace location</i>										
	Full model				Fixed floorspace				Fixed non-parents	
	Parents		Non-parents		Parents		Non-parents		Parents	
	High	Low	High	Low	High	Low	High	Low	High	Low
Log dist. center	0.016*** (0.002)	-0.002 (0.002)	-0.025*** (0.002)	-0.027*** (0.002)	0.032*** (0.001)	0.020*** (0.002)	-0.008*** (0.001)	-0.005*** (0.000)	0.002 (0.004)	-0.017*** (0.004)
Obs.	2,419	2,419	2,419	2,419	2,419	2,419	2,419	2,419	2,419	2,419
R ²	0.016	0.000	0.038	0.040	0.205	0.059	0.098	0.063	0.000	0.007

Notes: The unit of observation is the *sección*. Each column regresses the log change in residential share (Panel A) or workplace employment (Panel B), relative to the free-choice baseline, on log distance to the center (Plaza del Sol). Full model is the estimated barrio catchment policy; fixed floorspace and fixed non-parents hold, respectively, the residential-commercial floorspace split and non-parents' location choices at their baseline values. Workplace choice does not depend on parenthood directly (Section 4.3), so any parenthood gap in Panel B reflects the residential gap in Panel A. Non-parent columns under fixed non-parents are omitted, since that scenario holds non-parents' locations fixed by construction and those coefficients are mechanically zero. Heteroskedasticity-robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.7: The city-wide effects of catchment areas: households, workers, prices, and floorspace

Panel A: Wages and total workers									
	Full model			Fixed floorspace			Fixed non-parents		
	High wage	Low wage	Workers	High wage	Low wage	Workers	High wage	Low wage	Workers
Log dist. center	-0.005*** (0.001)	-0.005*** (0.001)	-0.020*** (0.002)	-0.000*** (0.000)	0.000*** (0.000)	0.002*** (0.000)	-0.008*** (0.001)	-0.008*** (0.001)	-0.031*** (0.004)
Obs.	2,419	2,419	2,419	2,419	2,419	2,419	2,419	2,419	2,419
R ²	0.029	0.031	0.022	0.010	0.019	0.086	0.027	0.032	0.022
Panel B: House prices and total households									
	Full model		Fixed floorspace		Fixed non-parents				
	Price	Households	Price	Households	Price	Households	Price	Households	
	Log dist. center	0.013*** (0.001)	0.033*** (0.005)	0.020*** (0.002)	0.021*** (0.002)	0.020*** (0.002)	0.048*** (0.006)		
Obs.		2,419	2,419	2,419	2,419	2,419	2,419		
R ²		0.041	0.032	0.045	0.046	0.040	0.036		
Panel C: Total, residential, and commercial floorspace									
	Full model			Fixed floorspace			Fixed non-parents		
	Total	Res.	Comm.	Total	Res.	Comm.	Total	Res.	Comm.
Log dist. center	0.007*** (0.001)	0.015*** (0.004)	-0.043*** (0.004)	-0.000 (0.000)	-0.000 (0.000)	-0.000* (0.000)	0.010*** (0.001)	0.021*** (0.005)	-0.065*** (0.006)
Obs.	2,419	2,419	2,419	2,419	2,419	2,419	2,419	2,419	2,419
R ²	0.032	0.013	0.034	0.002	0.001	0.003	0.032	0.016	0.033


Notes: The unit of observation is the *sección*. Each column regresses the log change in the variable, relative to the free-choice baseline, on log distance to the center (Plaza del Sol). Full model is the estimated barrio catchment policy; fixed floorspace and fixed non-parents hold, respectively, the residential-commercial floorspace split and non-parents' location choices at their baseline values. Heteroskedasticity-robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix B: Data construction

B.1 Travel times

I predict mode shares with a multinomial logit model of mode choice, using trip-level data from the EDM2018. The dependent variable is the reported mode (car, public transit, or walking) for each trip. I include travel distance by car as a regressor, capturing the pure length of the trip, along with the time differences between driving and each of the other two modes (transit and walking), which capture the relative attractiveness of each alternative for a given distance. Distances and travel times come from Google Maps, queried between the origin and destination *barrios* of each trip. Their centroids are obtained as the population-weighted average of the *secciones* that compose them, so that they do not fall in uninhabited areas.

Figure B.1: Predicted mode shares by distance bin, trip purpose, and skill level.



Notes: The graphs show the predicted share of each transport mode (car, transit, walking) across distance bins, used to construct weighted travel times throughout the paper. Panels (a) and (b) report shares for school trips, by high- and low-skill households respectively; Panels (c) and (d) report the same for work trips. Shares are estimated from the EDM2018 travel survey.

Denote by $m \in \{\text{car, transit, walk}\}$ the mode chosen for trip ω , with car as the base category. The multino-

mial logit model is:

$$\Pr(m_\omega = m) = \frac{\exp(\eta_m(\omega))}{\sum_{m' \in \{\text{car}, \text{transit}, \text{walk}\}} \exp(\eta_{m'}(\omega))}, \quad \eta_{\text{car}}(\omega) \equiv 0, \quad (\text{B.1})$$

where, for $m \in \{\text{transit}, \text{walk}\}$, the linear index is

$$\eta_m(\omega) = X'_\omega \beta_m, \quad X_\omega = (\text{dist}_\omega, \Delta_{\text{transit}, \omega}, \Delta_{\text{walk}, \omega}) \otimes (1, \text{skill}_\omega, \text{purpose}_\omega, \text{skill}_\omega \times \text{purpose}_\omega), \quad (\text{B.2})$$

with dist_ω the car travel distance for trip ω , and $\Delta_{m, \omega} \equiv \text{car_time}_\omega - \text{time}_{m, \omega}$ the time difference between driving and mode m . I estimate the model using survey weights, and predict mode shares for each distance, skill level, and trip purpose. These shares are then used to aggregate travel times across modes into a single value used throughout the paper.

For origin-destination pairs within the same geographic unit, where Google Maps returns no meaningful route, I instead impute travel time from the within-unit trips recorded in the EDM2018, averaging their reported duration by skill group.

Figure B.1 reports the results: mode shares vary with travel distance, skill, and trip purpose. For distances within 5 kilometers, for instance, walking is far more likely for school trips than for work trips. Low-skill households are also more likely to use public transit overall, so that, for the same origin-destination pair, their predicted travel time is higher than for high-skill households.

B.2 School admission data

The admission data from the Community of Madrid cover all students admitted to a public or concerted school who applied, at any level of compulsory schooling, from 2012 to 2015. For each student I observe the application school, the admission school, and other characteristics, including the residential coordinates, which I match to the corresponding *sección*.

I build a student-school pair dataset for each year to study the determinants of school choice. The main variables are the travel times, computed as described above, and the admission indicator. To construct the latter, I compute the admission points each student would obtain at every school. When a student's points exceed the observed cutoff for a school-year, the indicator equals one and the school enters the student's choice set. The cutoff is set to zero when the number of admitted students is below the available seats in that school-year. For oversubscribed schools, it equals the score of the lowest-admitted student.

The estimation sample applies several restrictions. I keep only students entering the first year of pre-school, the point of entry into the system and the only application most students make over their school career. I keep only students living in the city of Madrid, for whom I observe the application scores and can predict scores at other schools, since outside the municipality I lack the priority points granted by the proximity criteria. I keep only students admitted to their preferred school, who account for 92% of first-year pre-school applicants, since the others violate the full-information and sophistication assumptions.

I drop students whose fixed points fall outside the range permitted by the policy rules, likely a result of misreporting. I exclude students with sibling or alumni-family priority, whose choices depend on past conditions, at the time the sibling or relative was admitted, that may differ from those observed. I drop students who received proximity points through a parent’s workplace, since I do not observe where the parents work and cannot compute their counterfactual points at other schools. I drop students who filed a claim against the allocation, since I do not observe whether it changed their scores. When two or more students attend the same school with identical points and the same location, I treat them as twins and count their choice once. Finally, I remove students whose observed school lies outside their feasible set, which is due to misreporting in the original data.

Table B.1 reports statistics for the indicators of scoring points that determine the sample selection, among students enrolling in the first year of pre-school in the municipality of Madrid.

Table B.1: Prevalence of point-based exclusions among applicants

	Obs.	Share	Var.
Sibling priority	84,608	0.224	0.174
Alumni-family priority	84,608	0.063	0.059
Workplace application	84,608	0.067	0.062
Filed an appeal	84,608	0.011	0.011
Passes all point-based filters	84,608	0.670	0.221

Notes: Shares are computed among first-year pre-school applicants in the municipality of Madrid, pooling years 2012-2015. The criteria reported may overlap. The last row reports the share passing all four point-based exclusions jointly.

Appendix C: Model derivation details

C.1 Firms

Profit maximization under perfect competition gives factor demands

$$L_i = \frac{\beta Y_i}{W_i}, \quad H_i^C = \frac{(1-\beta)Y_i}{r_i^C}, \quad W_i = [\psi_i^\rho (w_{iH})^{1-\rho} + (1-\psi_i)^\rho (w_{iL})^{1-\rho}]^{\frac{1}{1-\rho}}, \quad (\text{C.1})$$

where W_i is the firm's unit cost of labor. Zero profits pin down the price of commercial floorspace:

$$r_i^C = \left(\frac{\beta^\beta (1-\beta)^{1-\beta} A_i}{W_i^\beta} \right)^{\frac{1}{1-\beta}}. \quad (\text{C.2})$$

C.2 Floorspace supply

Developers combine land T_n with capital K_n :

$$H_n = K_n^{\mu_n} (T_n \zeta_n)^{1-\mu_n}, \quad (\text{C.3})$$

where ζ_n is a location-specific land productivity fundamental, and $\mu_n \in \{\mu_C, \mu_S\}$ is the elasticity of floorspace to capital, allowed to differ between central and suburban locations.⁴¹ Developers evaluate the return to floorspace at the residential rent r_n , since residential and commercial rents are linked by an exogenous wedge, $r_n^C = \xi_n r_n$. The developer's first-order condition for capital gives

$$K_n = [\mu_n r_n (T_n \zeta_n)^{1-\mu_n}]^{\frac{1}{1-\mu_n}}.$$

Substituting into (C.3) gives a closed-form floorspace supply:

$$H_n = \mu_n^{\frac{\mu_n}{1-\mu_n}} \zeta_n T_n r_n^{\frac{\mu_n}{1-\mu_n}}. \quad (\text{C.4})$$

Residential and commercial floorspace draw on the same underlying stock H_n , so the residential share σ_n is not separately chosen: it equates to the relative expenditure on each use, once commercial expenditure is converted to residential-equivalent units through the wedge ξ_n :

$$\sigma_n = \frac{E_n^R}{E_n^R + E_n^C / \xi_n}, \quad (\text{C.5})$$

where $E_n^R \equiv r_n H_n^R$ is residential floorspace expenditure from equation (13), and $E_n^C \equiv r_n^C H_n^C$ is commercial floorspace expenditure from equation (C.1). Workplaces and neighborhoods index the same set of locations, so i and n denote the same location here. Residential and commercial supply are then $H_n^R = \sigma_n H_n$ and $H_n^C = (1 - \sigma_n) H_n$.

⁴¹Capital is supplied competitively in a single economy-wide market, at a price I normalize to one.

Appendix D: Additional estimation results

D.1 Workplace choice estimation

Table D.1 reports the travel-time elasticities under three departures from the baseline sample and specification. Separating the commuting matrices by skill and parenthood does not generate significantly different semi-elasticity estimates with both origin and destination fixed effects. Restricting the sample to *barrios*, and thus to the municipality of Madrid, yields slightly higher estimates. Setting travel time to zero for pairs sharing the same origin and destination, rather than using the average EDM2018 reported time as in the baseline, yields slightly lower estimates. Overall, the main results are confirmed.

Table D.1: Commuting gravity equation: robustness

	With children		Without children		Madrid only		Same-location zeros	
	High (1)	Low (2)	High (3)	Low (4)	High (5)	Low (6)	High (7)	Low (8)
Travel time	-0.079*** (0.003)	-0.098*** (0.003)	-0.085*** (0.003)	-0.095*** (0.002)	-0.094*** (0.004)	-0.105*** (0.005)	-0.075*** (0.002)	-0.083*** (0.001)
Origin FE	✓	✓	✓	✓	✓	✓	✓	✓
Destination FE	✓	✓	✓	✓	✓	✓	✓	✓
First-stage F	93785.1	115379.8	93785.1	115379.8	60599.6	76079.7	93733.0	114922.6
Observations	31,684	32,041	31,504	32,041	16,384	16,384	32,041	32,041

Notes: The dependent variable is the share of commuters from an origin to a given destination. Travel times are measured in minutes and differ between the high- and low-skill estimations, as each is a skill-specific weighted average across transport modes. All columns instrument travel time with straight-line distance through a control function, with standard errors bootstrapped over 100 replications. Columns (1)-(2) restrict the sample to households with children and columns (3)-(4) to those without; columns (5)-(6) keep only origin-destination pairs within the municipality of Madrid; columns (7)-(8) set travel time to zero for trips whose origin and destination coincide. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table D.2 reports the workplace choice parameters recovered via the gravity regression and variance matching procedures. The travel-cost elasticity $\kappa^W = 0.019$ is higher than the value typically used in the literature (around 0.01). The pattern is consistent with the semi-elasticities ϕ_g^W , which are also somewhat above previous estimates.

Table D.2: Workplace choice parameters

	ϕ_g^W (1)	θ_g^W (2)	κ^W (3)
High skill	0.082	4.366	0.019
Low skill	0.096	5.087	0.019

Notes: Column (1) reports the travel-time semi-elasticity ϕ_g^W from Table 1. Columns (2)-(3) report the dispersion parameter θ_g^W and the travel-cost elasticity κ^W . A common κ^W is imposed across skill groups.

D.2 Admission probabilities and skills

I treat the admission probability $\bar{p}_{nj}$ as common across skill groups within a *sección*. The assumption has an institutional basis, as the priority-point system rests mainly on characteristics unrelated to skill, with only a small component (points for low-income status) that could correlate with it (see Figure A.1). It is also a data constraint, since the application records contain no skill or education measure; the only skill information I observe is the share of high-skill households at the *sección* level. Measuring $\bar{p}_{nj}$ as a location average is therefore both consistent with the point system and the finest split the data allow.

At the *sección* level, admission probabilities barely covary with the local high-skill share. The share explains only about 4% of the within-year variation in average admission probability across *secciones*, and the association is small in magnitude: a one-standard-deviation increase in the high-skill share (0.27) raises the admission probability by 1.2 percentage points, roughly 0.17 standard deviations, against a mean of 0.85. Admission probabilities therefore barely covary with skill composition, which supports a common $\bar{p}_{nj}$.

D.3 Approximation of school access

In this section, I outline the reasons for and the consequences of the approximation I use to quantify school access in Equation 9. School access depends on the household ω through its admission probabilities $p_{\omega j}$. Carrying ω into the residential problem would require household-level information at the residence as well. Because I have no microdata linking school and residential choices, I aggregate school access to the location-by-household-type level. The aggregation introduces a bias in the measurement of school access, which I quantify next.

Write the household's inclusive value as $A_{\omega g} = \sum_j (p_{\omega j})^{\theta_g^s} u_{njg}$, where $u_{njg} = (S_{jg}/\tau_{njg}^s)^{\theta_g^s}$ is the utility term of school j for a household with skill g living in n . Under the sophistication and full information assumption introduced in the school choice problem, $p_{\omega j}$ is an indicator, so raising it to θ_g^s leaves it unchanged. Moreover, the admission term $p_{\omega j}$ is common across skill, as discussed in Appendix D.2. Finally, within a *sección* n , u_{njg} does not vary across households, since all face the same quality and travel costs. Thus, the school access approximation uses the following average inclusive value

$$\bar{A}_{ng} \equiv \frac{1}{N_n} \sum_{\omega \in n} A_{\omega g} = \sum_j \bar{p}_{nj} u_{njg}, \quad (\text{D.1})$$

which is constructed via the average local admission probability

$$\bar{p}_{nj} = \frac{\sum_{\omega \in n} p_{\omega j}}{\sum_{\omega} \mathbb{1}\{\omega \in n\}}. \quad (\text{D.2})$$

Exact school access in n averages the concave transform of $A_{\omega g}$ across households in n , while the approx-

imation evaluates the aggregator at the average option value:

$$S_{ng}^{exact} = \gamma_g^S \frac{1}{N_n} \sum_{\omega \in n} A_{\omega g}^{1/\theta_g^S} \leq \gamma_g^S \bar{A}_{ng}^{1/\theta_g^S} = \gamma_g^S \left(\sum_j \bar{p}_{nj} u_{njg} \right)^{1/\theta_g^S} = S_{ng}^{approx}. \quad (D.3)$$

Since $\theta_g^S > 1$, the aggregator x^{1/θ_g^S} is concave, so the approximation weakly overstates exact access via Jensen's inequality. A second-order expansion of x^{1/θ_g^S} around $\bar{A}_{ng} = \frac{1}{N_n} \sum_{\omega \in n} A_{\omega g}$ gives the magnitude of the gap:

$$S_{ng}^{approx} - S_{ng}^{exact} \approx \frac{\gamma_g^S}{2} \frac{1}{\theta_g^S} \left(1 - \frac{1}{\theta_g^S} \right) \bar{A}_{ng}^{1/\theta_g^S - 2} \text{Var}_{\omega \in n}(A_{\omega g}). \quad (D.4)$$

The bias from the second-order expansion is proportional to the within-sección variance of the option value $A_{\omega g}$, and vanishes when households in a sección face the same admission probabilities, so that $\text{Var}_{\omega \in n}(A_{\omega g}) = 0$. In what follows I report the exact difference between the two measures, rather than the second-order approximation.

Empirically, two features of the setting keep $\text{Var}_{\omega \in n}(A_{\omega g})$ low. First, secciones are small, so households within one share most of their geography and their nearby schools. Second, dropping applications with sibling priority removes the largest household-specific source of admission differences, leaving admission prospects close to common within a sección. Where admission is uncertain, it is typically at distant or oversubscribed schools, which carry little weight u_{njg} and therefore move $A_{\omega g}$ little.

Table D.3 reports the quantification exercise for the school-access approximation. The admission indicator has an unconditional variance of 0.126, falling to 0.029 within sección-school cells, so location means absorb 77% of its variation. Another way to gauge the bias is to count the cells with zero within variance, where the approximation is exact: the indicator is degenerate, all applicants facing the same admission probability, in 88.2% of sección-school cells, so the bias stems from a minority of cells. The relative bias across secciones has a median of 0.006% for high-skill and 0.002% for low-skill households, with a 90th percentile of at most 0.21%. The approximation is therefore negligible.

D.4 School choice estimation

Table D.4 summarizes the 2015 estimation sample for school choice. Two features are worth noting, as they shape the application context and behavior. First, admission is nearly universal. The average applicant clears the cutoff at 451 of the 514 schools in the city, and the admission indicator averages 0.88 across all student-school pairs. The feasible set is therefore large, and admission screens few schools, so the distribution of travel times to feasible schools is virtually identical to that over all schools. Second, applicants do not choose the closest option but trade proximity against other features. They attend a school 7.8 minutes away on average, more than twice the 3.5 minutes to their closest school. The pattern lets me quantify the travel-time sensitivity and recover perceived school quality from revealed preference.

Figure D.1 offers another way to see that travel times matter but cannot fully explain school choices. It plots

Table D.3: School-access approximation: admission probabilities and induced bias

	Unit	Obs.	Mean	Var.	p10	p50	p90
Admission indicator	student-school	20,052,423	0.853	0.126	0	1	1
residualized on <i>sección</i>	student-school	20,052,423	-0.000	0.029	0	0	0
Degeneracy	<i>sección</i> -school cell	4,600,519	0.882	0.104	0	1	1
Bias, high skill (%)	<i>sección</i>	9,061	0.094	0.104	0	0.006	0.211
Bias, low skill (%)	<i>sección</i>	9,061	0.086	0.131	0	0.002	0.175

Notes: Statistics in the table pool years 2012-2015. The admission indicator equals one when a student clears a school's cutoff. The second row residualizes it on *sección*-school-year means, so its mean is zero by construction and its variance is the within-cell dispersion that matters for the approximation. Degeneracy is an indicator that a *sección*-school-year cell has no within variation (all admitted or all rejected). The bias rows report $100 \times (\mathbb{S}_{ng}^{\text{approx}} - \mathbb{S}_{ng}^{\text{exact}}) / \mathbb{S}_{ng}^{\text{exact}}$ across *secciones*, where $\mathbb{S}^{\text{exact}}$ averages household-level access and $\mathbb{S}^{\text{approx}}$ uses the *sección*-average admission probability.

Table D.4: School choice: descriptive statistics

	Unit	Obs.	Mean	Var.	p10	p50	p90
High-skill probability	student	9,033	0.57	0.07	0.20	0.58	0.92
Number of feasible schools	student	9,033	451.01	1390.66	409.00	477.00	479.00
Travel time to chosen school	student	9,033	7.84	27.80	2.44	6.59	15.05
Travel time to closest school	student	9,033	3.53	5.28	1.37	3.10	5.85
Admission indicator	pair	4,642,962	0.88	0.11	0.00	1.00	1.00
Travel time to feasible schools	pair	4,073,980	21.10	27.92	15.16	20.73	28.01
Travel time to all schools	pair	4,642,962	21.13	27.50	15.23	20.74	28.03


Notes: Sample of first-year pre-school applicants resident in Madrid for 2015. Travel times for high-skill measure are in minutes; low-skill travel times are slightly higher but highly correlated. The number of feasible schools counts schools a student clears on admission. Travel time to the closest school is the minimum over all schools, ignoring admission. Units are either individual students or a student-school pair.

the cumulative share of students by the distance rank of their chosen school from their residential *sección*. Already 22% of students choose their closest option, and the fourth-ranked school brings the cumulative share over half. Beyond that, successive schools add less and less, with the tenth school only reaching 75% of the cumulative share. The evidence again suggests that distance matters, but that many families do not choose on distance alone, weighing something else, bundled in the school-quality term that I recover from revealed preference.

D.5 Estimation of expenditure share by skill and parenthood

Table D.5 reports the housing share of income by parenthood and skill. I use the Encuesta de Presupuestos Familiares (EPF) from the INE for the Community of Madrid to estimate how parenthood and skill affect the share of income a household spends on housing, controlling for the age of the principal earner, the number of income earners, the survey year, municipality size, and local density. Skill shifts the share substantially: low-skill households devote about 44% of income to housing against 37% for high-skill households, a gap of roughly seven percentage points that is precisely estimated and robust to controls. Parenthood, by con-

Figure D.1: School choice by distance rank



Notes: The figure reports the share of students choosing school by their rank distance order within their feasible set, for the final dataset in 2015.

trast, leaves the share unchanged, when adding controls. The *Family* coefficient is small and statistically insignificant once controls are added, and its interaction with skill is insignificant throughout, so the absence of a parenthood effect holds for both skill groups. Housing shares therefore depend on skill but not on parenthood, which justifies collapsing α_{gk} to α_g and lets the parenthood comparison difference out house prices in Equation 23.

D.6 Peer-effects estimation

I assemble a school-level dataset to study peer effects. I focus on the cross-section of schools in the municipality of Madrid in 2015 that offer the first year of pre-school. After the cleaning applied to the school-choice sample, the data include 514 schools with non-missing estimated quality and information on admitted students. The exact location of each school lets me measure its elevation from a terrain model, its distance from the center (*Puerta del Sol*), and its construction year from the Spanish cadastre. Other characteristics are scraped from the school-choice information system of the Madrid regional government (*Buscador de colegios*) and from an online forum (*Micole.net*). The average school has more students from high-skill backgrounds and a little over 40 students in total. Half are concerted, and 40% have a religious affiliation. More than 90% carry no distinct price tier, most teach both Spanish and English, and only a few offer alternative pedagogy or additional languages. Relevant for identification, the p_{10} - p_{90} range of elevation is narrow, roughly 100 meters, too small to drive systematic differences in fundamental school quality.

Figure D.2, Panel B, shows the average elevation for each *sección*, with lighter areas indicating lower elevation and tracing the path of the Manzanares, which runs south of the historical center. Panel A shows that elevation correlates strongly with the ratio of high- to low-skill households. The range of elevation within the municipality is narrow, less than 200 meters between peak and trough. Elevation is likely therefore less salient to households today, though it was historically meaningful when wastewater flowed south toward

Table D.5: Housing expenditure share by parenthood and skill

	Housing share of income	
	(1)	(2)
Family	-0.095*** (0.008)	-0.005 (0.009)
High-Skill	-0.105*** (0.010)	-0.066*** (0.009)
Family $\times$ High-Skill	0.010 (0.012)	-0.006 (0.011)
Controls	No	Yes
Obs.	8,716	8,715
R2	0.057	0.234

Notes: Sample of households in the Community of Madrid, EPF 2010–2015. The outcome is housing expenditure as a share of income. *Family* indicates the presence of dependent children; *High-Skill* indicates a college-educated principal earner. Column (2) controls for age, survey year, municipality size, density, and number of income earners. Heteroskedasticity-robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.


Table D.6: School-level descriptive statistics

	Obs.	Mean	SD	p10	p50	p90
Quality, high-skill	514	1.01	0.15	0.81	1.01	1.20
Quality, low-skill	514	1.01	0.16	0.81	1.01	1.20
Skill ratio (H/L)	514	1.51	1.58	0.27	0.92	3.70
Admitted students	514	43.21	29.87	10.00	39.00	83.00
Elevation (m)	514	654.33	40.59	593.30	662.69	704.14
Construction year	444	1970.76	23.65	1946.00	1970.00	2004.00
Distance to center (km)	514	8.38	3.84	3.81	7.72	13.61
Concerted	514	0.52	0.50	0.00	1.00	1.00
Religious	501	0.40	0.49	0.00	0.00	1.00
High price	501	0.00	0.00	0.00	0.00	0.00
Med.-high price	501	0.01	0.09	0.00	0.00	0.00
Med.-low price	501	0.08	0.27	0.00	0.00	0.00
Low price	501	0.91	0.29	1.00	1.00	1.00
Alternative pedagogy	501	0.28	0.45	0.00	0.00	1.00
Spanish-English	501	0.83	0.38	0.00	1.00	1.00
More languages	501	0.13	0.34	0.00	0.00	1.00

Notes: The unit of observation is the school, for the full cross-section used in the peer-effects estimation. Quality is the log school quality recovered from the school-choice estimation. The skill ratio is the number of high-skill over low-skill admitted students. Rows below the second break are indicators, so their means are shares. Elevation, construction year, and distance are in levels; the estimation uses their logs.

the river. The narrow range supports the exclusion restriction, since elevation differences are too small to enter school choice today except through the historical sorting they once produced.

Figure D.2: Elevation above sea level



Notes: Panel a: the scatter and fit between *sección* elevation (x-axis) and the ratio of high- to low-skill households (y-axis). Panel b: the elevation of each *sección* in the municipality of Madrid, with darker colors indicating higher elevation. In the south, the lighter units trace the Manzanares.

Finally, I report the first stage of the peer-effects estimation. Table D.7 gives the effect of elevation on the number of admitted students by skill background and on the high-to-low-skill ratio. Elevation strongly predicts more admissions of high-skill students and fewer of low-skill students. The relationship follows because elevation predicts the high-skill share at the *sección* level, household skill is inferred from residence, and school choice is spatially concentrated near the residence, as the school-choice estimation shows. The elasticity of the skill ratio to elevation is large, at 8.5, so schools in higher locations admit disproportionately more high-skill students, conditional on the controls.

D.7 Residential choice estimation

To estimate residential preferences I build a dataset at the *sección* level, which includes the likelihood of each group of households to live there from the Padrón Municipal, the constructed school access, the average school construction year instrument and other relevant controls, such as the distance to the city center, the elevation above sea level, and the characteristics of the housing stock, such as the average construction year and average dwelling size. Table D.8 presents the descriptive statistics of the main variables used in the residential preference estimation. The higher ratio for high-skill households indicates that families with children are relatively more concentrated among the high-skill population, consistent with stronger residential sorting by high-skill parents. School access is similar across skill groups, with a somewhat contained dispersion. Importantly, the construction years point to a gap between schools and housing. Schools around each *sección* appear to be somewhat older than the housing within them. This is what allows the instrument to exploit school-building timing separately from neighborhood housing vintage.

Table D.7: Student skill ratio and elevation

	All schools			Large schools		
	High (1)	Low (2)	Ratio (3)	High (4)	Low (5)	Ratio (6)
Log elevation	5.254*** (0.779)	-3.205*** (0.518)	8.459*** (0.820)	5.320*** (0.766)	-3.218*** (0.520)	8.538*** (0.841)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Size				≥ 15	≥ 15	≥ 15
Obs.	441	441	441	369	369	369

Notes: The unit of observation is the school. The outcome is the log number of high-skill students, low-skill students, or their log ratio. Columns (4)-(6) restrict to schools with at least 15 admitted students. All specifications weight by total admitted students and control for construction year, distance to the center, school ownership, religious status, price tier, and pedagogical and language offerings. Heteroskedasticity-robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table D.8: Sección-level descriptive statistics

	Obs.	Mean	SD	p10	p50	p90
Residential share ratio, high-skill	2,420	1.06	0.71	0.50	0.87	1.77
Residential share ratio, low-skill	2,420	0.95	0.64	0.37	0.81	1.60
School access, high-skill	2,245	1.11	0.09	1.01	1.11	1.21
School access, low-skill	2,245	1.11	0.10	0.98	1.11	1.23
Nearby school construction year	2,395	1966.11	16.79	1945.55	1969.01	1981.66
Building construction year	2,420	1970.67	15.92	1952.00	1970.00	1993.00
Dwelling size (m ²)	2,420	84.86	29.67	60.26	77.87	114.01
Elevation (m)	2,420	655.98	39.74	597.01	665.04	704.02
Distance to center (km)	2,420	7.06	3.03	3.27	6.68	10.82

Notes: The unit of observation is the residential *sección*. The average school construction year considers schools within 1.5km from the *sección*, and weights their importance by inverse distance.


Next, I show evidence on the first stage in Table D.9. The average construction year of nearby schools strongly predicts school access for both skill groups, with positive and significant coefficients across all specifications, confirming the instrument relevance. The relationship holds when *barrio* fixed effects are added, though the doubling of the standard errors reflects the loss of identifying variation absorbed by the local fixed effects. Another way to observe the relevance of the instrumental variable is to switch to the school level and check the relationship between school quality and the construction year of the building. Figure D.3, Panel A, shows the scatter for both high and low-skill school quality, and both are shown to be increasing in the year of construction, likely because parents prefer newer buildings. As expected, the map in Panel B, shows how central areas in Madrid have access to schools that are on average older.

Table D.9: First stage: school access and school construction year

	High-skill access		Low-skill access	
	(1)	(2)	(3)	(4)
Log avg. school constr. year	1.694*** (0.374)	2.160*** (0.675)	2.970*** (0.470)	2.408*** (0.840)
Controls	Yes	Yes	Yes	Yes
Barrio FE		Yes		Yes
Obs.	2,222	2,218	2,222	2,218

Notes: The unit of observation is the residential *sección*. The outcome is log school access and the regressor is the log average construction year of schools within 1.5km, weighted with the inverse distance. Controls are the average dwelling size and building age of the *sección*, elevation and distance to the center. Columns (2) and (4) add *Barrio* fixed effects. Heteroskedasticity-robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure D.3: School construction year



Notes: Panel a: the scatter and fit between school construction year (x-axis) and school quality (y-axis). Panel b: the average school construction year within 1.5km from each *sección*, weighted by the inverse distance.

Finally, I run robustness analysis for the residential-preferences estimation. Table D.10 reports the pre-

ferred IV specification under changes to the dependent variable and the instrument. The effect of school access is positive and significant in every specification, for both skill groups. Redefining a family as a household with a child under five, against households whose youngest member is at least 34, leaves the estimates close to the baseline, so the result does not hinge on the age cutoff. The estimates are more sensitive to the instrument, particularly for the high-skill group. The smallest coefficients come from the 500-meter radius, which has the strongest first stage at the potential cost of a significant sample selection bias, while the largest come from the access-term instrument, which has the weakest. Read together, the exercise confirms the sign and significance of the residential response to school access, and its magnitude in the better-identified specifications.

Table D.10: Residential preferences: robustness

<i>Panel A: High-skill households</i>					
	(1) Alt. family	(2) IV radius 500m	(3) IV radius 3km	(4) IV without weights	(5) IV access term
Log school access	8.684*** (2.215)	4.459*** (1.319)	16.112*** (5.899)	13.784*** (4.730)	32.123* (19.226)
F-stat	20.39	25.68	6.65	9.43	2.77
Obs.	2,221	1,609	2,238	2,222	2,244
<i>Panel B: Low-skill households</i>					
	(1) Alt. family	(2) IV radius 500m	(3) IV radius 3km	(4) IV without weights	(5) IV access term
Log school access	6.713*** (1.573)	3.608*** (1.100)	9.846*** (2.366)	6.610*** (1.496)	13.405*** (3.438)
F-stat	39.04	40.65	19.53	34.97	18.43
Obs.	2,207	1,609	2,238	2,222	2,244

Notes: The unit of observation is the residential *sección*, and all columns instrument log school access. Column (1) redefines families as households with a child under five relative to those whose youngest member is at least 34. Columns (2) and (3) vary the instrument radius to 500 meters and 3 kilometers. Column (4) uses the unweighted average school construction year within 1.5 km. Column (5) uses the construction-year instrument built with the same structure of school access. All columns include the baseline controls. Heteroskedasticity-robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D.8 Inversion and recovery of fundamentals

The model fundamentals are recovered recursively, in the order stated in Section 5.5. Each step takes the observed data from Madrid and the previously recovered objects as given:

1. **Skill intensity.** Given the skill-specific labor allocations implied by the residential data and the labor supply equation, and given recovered wages, the labor first-order condition from the production function (Equation 15) yields $\psi_i = [1 + (L_{iL}/L_{iH})^{1/\rho} (w_{iL}/w_{iH})]^{-1}$, the skill-intensity in every *sección* of the city.
2. **Floorspace prices and wedge.** Residential prices at the *sección* level come from observed rents in tax declarations. I then aggregate residential expenditure by each group, given prices, incomes and residential shares. Unobserved commercial prices follow from firm expenditure (obtained from aggregate labor and wages, given the Cobb-Douglas production function) and observed commercial floorspace. The floorspace allocation wedge ξ_n is then given by the ratio of commercial to residential prices.
3. **Productivity.** Total factor productivity follows from the previously recovered wages and commercial floorspace prices, and the fundamental v_i^A nets out agglomeration according to Equation 16.
4. **Land productivity.** ζ_n inverts the floorspace supply technology in Equation 17, given total built-up floorspace, prices, and the center or periphery supply elasticities. I define the dense center of the city for the lower floorspace elasticity as the areas that were already built up in the 1960s.
5. **School quality.** v_{jg}^S inverts Equation 19, netting the recovered quality of the peer-effects term, with the student composition implied by the residential data via school choice.
6. **Amenities.** B_{ngk} inverts the residential-choice probabilities from Equation 12, as the residual reconciling observed residential shares with workplace access, school access, and prices. School access enters only for households with children. Amenities are normalized to a unit geometric mean.

D.9 Access decomposition into extensive and intensive margins

Consider any market access measure defined as the ex-ante expected utility from a discrete choice problem with Fréchet shocks. Any such measure splits exactly into an extensive and an intensive margin. The decomposition follows from previous literature, and builds on two ideas from other fields. The first is the inverse Herfindahl index from the IO tradition, which measures the number of equally important options (Adelman, 1969). The second is the split, drawn from trade and spatial economics, of a price index into the prices of available goods and their availability (Feenstra, 1994; Handbury and Weinstein, 2015).

I next illustrate the decomposition for a general market access term. To start, consider a household in n choosing among options indexed by j , each with a utility composite term A_{nj} , so that access and choice probabilities are respectively

$$A_n = \left(\sum_j A_{nj} \right)^{1/\theta}, \quad \lambda_{nj} = \frac{A_{nj}}{\sum_\ell A_{n\ell}}.$$

Following a long tradition in IO (Adelman, 1969), define the effective number of options as the inverse Herfindahl index of the choice probabilities, $N_n^{eff} = 1/\sum_j \lambda_{nj}^2$, and the average utility across the options

a household chooses among as $\bar{A}_n = \sum_j \lambda_{nj} A_{nj}$. Substituting for the choice probabilities λ_{nj} in both, and rewriting $\sum_j A_{nj} = A_n^\theta$, the two objects share the same sum of squared utility terms,

$$N_n^{eff} = \frac{A_n^{2\theta}}{\sum_j A_{nj}^2}, \quad \bar{A}_n = \frac{\sum_j A_{nj}^2}{A_n^\theta}.$$

Indeed, the product of the two leaves A_n^θ , so that

$$A_n = \underbrace{(N_n^{eff})^{1/\theta}}_{\text{extensive margin}} \times \underbrace{\bar{A}_n^{1/\theta}}_{\text{intensive margin}}, \quad \text{or} \quad \ln A_n = \frac{1}{\theta} \ln N_n^{eff} + \frac{1}{\theta} \ln \bar{A}_n.$$

To interpret the decomposition, one has to focus on the two parameters that govern access in quantitative urban models. The utility composite term typically depends on κ and θ ,

$$A_{nj} = (x_j / \tau_{nj})^\theta \quad \text{with} \quad \tau_{nj} = e^{\kappa t_{nj}},$$

where x_j is the value of option j and t_{nj} the travel time to reach it.

The dispersion parameter θ shapes the choice probabilities themselves. As θ grows, idiosyncratic shocks vanish and the probabilities concentrate on the options with the highest value, so that, in the limit, $N_n^{eff} \rightarrow 1$ and access collapses to the intensive margin. As θ falls, the probabilities spread across the reachable options and the extensive margin gains weight. The semi-elasticity of travel costs κ operates instead by discounting the utility of each option. Since the utility decays at rate $\theta\kappa$ in travel time, a higher κ drives the probabilities of distant options toward zero, which lowers N_n^{eff} and ties it to the available options close by. The balance between the two margins in space consequently depends on κ through the geographic concentration of options, and on θ through the weight that their value carries.

The decomposition is an identity, and it holds for any market access measure. Because ordinary least squares is linear in the dependent variable, the elasticity of access to any covariate splits into the elasticities of the two margins, $\beta^A = \frac{1}{\theta} \beta^N + \frac{1}{\theta} \beta^{\bar{A}}$, which lets me read the gradients of school and workplace access with distance to the center as the sum of an extensive and an intensive component (see Table 5).