

Yale University

EliScholar – A Digital Platform for Scholarly Publishing at Yale

Discussion Papers

Economic Growth Center

Fall 11-2-2024

The Welfare Effects of including Household Preferences in School Assignment Systems: Evidence from Ecuador

Gregory Elacqua
Inter-American Development Bank

Isabel Jacas
Econsult

Thomas Krussig
Econsult

Carolina Méndez
Inter-American Development Bank

Christopher Neilson
Yale University, christopher.neilson@yale.edu

Follow this and additional works at: <https://elischolar.library.yale.edu/egcenter-discussion-paper-series>



Part of the [Behavioral Economics Commons](#), and the [Growth and Development Commons](#)

Recommended Citation

G. Elacqua, I. Jacas, T. Krussig, C. Méndez, C.A. Neilson "The Welfare Effects of including Household Preferences in School Assignment Systems: Evidence from Ecuador", EGC Discussion Papers. 1110.

This Discussion Paper is brought to you for free and open access by the Economic Growth Center at EliScholar – A Digital Platform for Scholarly Publishing at Yale. It has been accepted for inclusion in Discussion Papers by an authorized administrator of EliScholar – A Digital Platform for Scholarly Publishing at Yale. For more information, please contact elischolar@yale.edu.

The Welfare Effects of including Household Preferences in School Assignment Systems: Evidence from Ecuador*

G. Elacqua I. Jacas T. Krussig C. Méndez C.A. Neilson

November 2, 2024

Abstract

We partner with Ecuador’s government to implement a centralized school choice system using a Deferred Acceptance algorithm in Manta. Our study evaluates the welfare impact of transitioning from a distance-based assignment system to one that incorporates families’ preferences. Results show that accounting for preferences yields substantial welfare gains. Counterfactual analyses suggest that alternative mechanisms offer limited improvements compared to the benefits of preference inclusion and coordinated assignments. Household survey data on beliefs and satisfaction support these findings, indicating that centralized school choice systems can deliver significant welfare effects in developing countries.

JEL: I20, I21, I22

Key words: Market Design, centralized student assignment, school choice, Ecuador

*This research paper is the result of a collaboration between the Ecuadorian Ministry of Education, the Inter-American Development Bank, and the implementing partner tech NGO Consiliumbots. The views and opinions expressed in this paper are those of the authors and do not necessarily reflect the views of the Inter-American Development Bank, its Board of Directors, or the countries they represent. We appreciate the collaboration from Ecuadorian Ministry of Education throughout this project. We thank the implementation team of data scientists and developers at ConsiliumBots for their help throughout the project. Gregory Elacqua and Carolina Méndez at the Inter-American Development Bank, Christopher A. Neilson at Yale University and Isabel Jacas and Thomas Krussig formerly at ConsiliumBots and now at Econsult. Coresponding author: christopher.neilson@yale.edu

1 Introduction

This paper examines the welfare effects of a policy change in Manta, Ecuador’s coordinated school assignment system, which previously aimed to minimize travel distance. The system faced challenges due to the need for accurate georeferencing and adjustments for geographic barriers like hills and rivers. To address these issues, the Ecuadorian Ministry of Education, in partnership with the IADB and ConsiliumBots, piloted a new system driven by applicants’ preferences. The new policy adopted best practices, including the deferred acceptance algorithm, unlimited ranked lists, and improved information systems (Abdulkadiroğlu & Sönmez, 2003; Pathak, 2011, 2017; Abdulkadiroğlu et al., 2017; Arteaga et al., 2021).

To compare the assignment alternatives, we take advantage of the fact that the system implemented in Manta elicited the true preferences and locations of all participating applicants. This allows us to compare the assignments made by the new centralized choice and assignment system (CCAS) with the simulated assignments of the prior alternative. Our methodology consists of using a counterfactual strategy, replicating the rules of the previous process, and simulating assignments with different lotteries following Abdulkadiroğlu et al. (2017).

Our key finding is that implementing a coordinated mechanism that accounts for applicants’ preferences provides significant welfare benefits. Using the deferred acceptance (DA) algorithm increased the percentage of applicants assigned to one of their chosen schools from 49.96% to 78.44%, and first-choice assignments rose from 42.42% to 69.76%. However, this comes with an average increase in distance to school by 0.29 km, with Preschool 1 and 2 applicants seeing increases of 0.683 km and 0.354 km, respectively, while Primary 1 applicants saw a smaller increase of 0.012 km due to greater congestion. For those whose assignments differ between the distance and DA mechanisms, these distance increases are larger.

We focus our analysis on estimated welfare and on the share of applicants being assigned to more or less preferred alternatives based on their reported preferences as detailed below. School quality is not considered because: i) we do not aim to study whether families in Ecuador prefer higher-quality schools, but rather to assess the welfare consequences of the assignment system as it relates to applicants’ valuation of different schools, and ii) we cannot (at least directly) observe school quality.¹

These results enhance the understanding of the benefits of coordinated school choice and assignment systems. While several studies highlight the welfare advantages of different mechanisms, few have directly estimated the benefits of coordinated systems that incorporate household preferences. The most closely related study is Abdulkadiroğlu et al. (2017), which examines the welfare effects of shifting New York City’s uncoordinated assignment system to a coordinated one that factors in family preferences. The authors find that most welfare gains come from coordination using the deferred acceptance algorithm, with only marginal improvements from alternatives. Similarly, in Ecuador, we find that the DA algorithm outperforms a coordinated system based on minimizing distance.

¹The Ecuadorian government does not currently apply census-based student learning assessments in primary grades. We also do not study the impact of the system on other measures of interest, such as educational segregation, as we lack socioeconomic data for participating applicants.

2 Context and Algorithm Descriptions

We study school assignment in the coastal region of Manta, Ecuador. Specifically, we concentrate on the urban areas within and around the city of Manta, including the geographic units (“*cantones*”) of Manta, Montecristi, and Jaramijó.

Manta was selected as the result of a process that aimed to find a small but representative city to scale up the school assignment policy.² The selection process took into account students in the urban area, school coverage, distribution of school types (mainly public and private), as well as city size. Ultimately, Manta was chosen for its relative similarity to the alternatives. Table 4 of Appendix B compares the main characteristics of Manta and Guayaquil, another coastal city and the country’s largest, using data from the 2010 Census and school transfer requests in the 2019-2020 school year.

The Ecuadorian educational system is organized into three levels: Preschool (*Educación Inicial*), Primary School (*Educación General Básica*) and Secondary School (*Bachillerato*). In this paper, we focus on school assignments at the “entry level”, a designation that encompasses enrollment in Preschool 1, Preschool 2, and Primary 1.

Ecuador has three types of schools: public (*fiscal* and municipal), *fiscomisional* (partially state-funded), and private (fully family-funded). Nationwide, public schools account for 73.8% of enrollments, with *fiscomisional* at 6% and private schools at 20%. In Manta, public schools represent 66% of entry-level enrollments, while *fiscomisional* and private schools account for 4% and 30%, respectively.

This pilot focused on free public schools, so private schools were an outside option not included in our model, potentially leading to overestimated welfare comparisons. However, a survey conducted after the application period showed limited overlap between public and private schools.³ Only 1% of respondents cited private schools as a reason for not listing more alternatives, despite many applicants submitting just one or two preferences.⁴ This suggests that private school options were not a major factor for most families.

2.1 Distance-Centric Algorithm

Before the COVID-19 pandemic, the school assignment system in Ecuador was based mainly on the applicant’s location, reported through the code on the family’s electricity account (CUEN). In addition to the linear distance criterion, a prioritization criterion was also used to determine the order in which applicants were processed (being processed first was preferable). This prioritization criterion was randomly assigned to students applying to be enrolled in the system.

²There was a change in government in Ecuador in 2021, and the new administration recently decided to scale up the system in coastal districts beginning in 2023. Because of the COVID pandemic and the fact that the distance-centric alternative required several in-person interactions during the process, the Ministry is currently using a First-Come, First-Serve digital system.

³The objectives of this survey were to gather information about parents’ overall satisfaction with the system, the information sources they used to apply for schools, awareness of the school supply, among other aspects. The survey was completed by 1,484 parents.

⁴All applicants were ultimately assigned, though some were placed in a school outside their reported preferences. In these cases, the assigned school was the closest possible alternative, as explained in subsection 2.1.

The assignment system was part of a broader six-phase enrollment process, detailed in Appendix C. The Assignment Phase, the third of these phases, was conducted in stages. In the first stage, applicants were categorized based on their preference for regular, rural, bilingual, or special education programs. Where applicable, applicants with siblings already enrolled were assigned to the same school. Students applying for non-regular programs and those with siblings in the system were prioritized and processed before other applicants.

Once these groups were assigned via a process that was carried out directly at the district headquarters, the rest of the students were assigned using the distance-centric (DC) algorithm, which the Ministry called the “mathematical model.” Legal guardians could complete an individual registration (of a single applicant) or one for a “group of siblings.” While this latter option suggests that the system prioritized assigning groups of siblings to the same school over distance-based considerations, this was not confirmed by the Ministry experts with whom we interacted.

The processing of regular assignments was carried out as follows:

- At each level, random numbers were given to all applicants. These random numbers correspond to the prioritization criteria mentioned above and determined the processing order.
- Following this order, applicants were assigned to the closest school (linear distance) with vacancies, in an iterative process that used increasing distance radii from the applicant’s home.⁵

This procedure can be conceptualized as an application of the “serial dictatorship” mechanism, in the sense that applicants select schools one after the other. Given that the order of choice has a random component, such assignment models have been termed “random serial dictatorships” (Abdulkadiroğlu & Sönmez, 1998). The enrollment of groups of siblings can thus be considered a priority, since the system will try to assign these groups to the same schools over other individual applications.

As explained above, an applicant’s home address was based on the legal guardians’ electricity bill. Using the latter to identify family location has proven highly effective, but may also incentivize families to procure (and even buy) electricity bills closer to their schools of interest. Moreover, there are still areas where households do not have electricity meters. These facts were reported in a series of interviews carried out by the IADB in Quito and Guayaquil, where families and officials recounted different factors affecting the registration processes.⁶ Given that we do not have precise estimates of location misreporting rates, we conduct a sensitivity analysis in Section 4 and simulate assignments under different levels of misreporting.

The Ecuadorian government’s concern with minimizing the distance to school arises from public policy considerations, and not because this aspect affects other dimensions such as, for

⁵The schools available were evaluated at radii of 100m, 200m, 300m, and 500m, and then at increments of 250m up to 3.5 kilometers.

⁶For example, district officials commented: (1) “In District 24, Durán, Guayaquil, families lend their electricity bills to each other so they can all have access to the education system. We estimate that more than 60% of families in this district do not use their own electricity bill, so they do not register their real geolocation.” (2) “In District 8, Monte Sináí, Guayaquil, families maintain that there are “illegal invasions” of other families in areas where popular schools are located, using electricity bills from that area to get a seat in these schools.”

example, public expenditure on free busing to schools. The latter consideration is nevertheless relevant in other contexts (e.g., many US cities), meaning that analyses comparing assignment mechanisms in similar cases should consider inclusion of these budget factors.

2.2 Deferred Acceptance (DA) Mechanism

The pilot used the deferred acceptance mechanism (Gale & Shapley, 1962), following the best practices in school choice mechanism design (Pathak, 2011; Correa et al., 2019). The specification of the assignment algorithm included static and dynamic sibling priorities, family linking, and a multiple tie-breaking rule.⁷

The static and dynamic sibling priorities indicate that an applicant will be prioritized for assignment to a school/program if their sibling is already assigned to the school (static). If the applicant is applying at the same time with another sibling, and one of them is assigned to a school,⁸ the applicant that has not been assigned yet will receive priority for being assigned to that same school (dynamic). The dynamic sibling priority is lower than the static sibling priority because the latter is already defined (the sibling is attending the school), while the former will depend on the answer from the applicant after the assignment.

The family linking feature consists of trying to assign all siblings applying together to the same schools. Following a descending order, where older applicants are assigned first, if an older sibling is assigned to school A, the applications of the younger siblings will be modified to put school A as the first-ranked school to improve the probability of being assigned together. Finally, a multiple tie-breaking rule gives each applicant a different lottery number for each school to which they apply. Lottery numbers are used to break ties within priority groups when a school receives more applications than spaces available.

3 Data

The data used in this paper come from the centralized choice and assignment system (CCAS) pilot web page created in 2021 in the region of Manta, Ecuador.⁹ The first data set comprises the supply of vacancies for all schools and programs offered in the pilot, where an educational program consists of a combination of grade and school. The pilot was implemented for all students entering Preschool 1, Preschool 2, and the first year of primary school (i.e., ages 3 to 5) for the first time. Vacancies are presented in Panel A of Table 1.

Preschool 1 has the most vacancies and is the least congested grade, while Primary 1 is the

⁷The deferred acceptance algorithm was chosen for its non-strategic and stable nature, and its ability to support dynamic sibling priority, family-linked applications, and varied priority-quota combinations. While more efficient alternatives like SIC and TTC exist, they sacrifice strategy-proofness and stability, respectively. However, as discussed in Section 4 and by Abdulkadiroğlu et al. (2017), the efficiency improvements from these alternatives are marginal compared to transitioning from an uncoordinated or non-CCAS system, as is relevant here.

⁸This can happen if one sibling is older than the other and will depend on the order in which the algorithm is run. If it is descending, the older sibling will give dynamic priority to the younger sibling. If it is ascending, it will be the other way around.

⁹All PII data was eliminated for that purpose.

most congested.¹⁰

The second data set contains information about students and their legal guardians, including geolocation, sibling relationships, special educational needs, and nationality.¹¹ Each applicant’s rank-ordered list (ROL) of preferences had no length limit, and lotteries were conducted for each program. As noted, applications were supplemented by adding all other alternatives, sorted by distance, to the initial preference list in case the applicant was not assigned to a preferred school. While each program listed by the student had a different lottery number, the same lottery number was used for the appended list of alternatives.

Table 1: Vacancies and Applicants by Geographic Unit (*Cantón*) and Grade

Panel A: Vacancies			
Cantón	Preschool 1	Preschool 2	Primary 1
Manta	1,830	1,394	425
Montecristi	905	668	654
Jaramijó	110	47	37
Total Grade	2,845	2,109	1,116
Total Global	6,070		

Panel B: Applicants			
Cantón	Preschool 1	Preschool 2	Primary 1
Manta	1,101	1,143	338
Montecristi	481	437	124
Jaramijó	125	125	107
Other	2	0	1
Total Grade	1,709	1,705	570
Total Global	3,984		

The distribution of applicants by geographic unit and grade is presented in Panel B of table 1. Notably, at least in the case of the geographic unit of Manta (*Cantón*), the number of applicants in Preschool 1 and Preschool 2 is roughly equivalent. Although this poses a challenge from a public policy standpoint in that it is desirable to enroll students earlier, it is also an interesting dynamic for the application system since families’ decision to postpone the enrollment of their child(ren) puts them at a strategic disadvantage. This is because there are fewer available seats in Preschool 2, given that currently enrolled Preschool 1 students move automatically to the next level.

Figure 2 of Appendix A provides an overview of applicant priorities and the lengths of ranked ordered lists. Note that most applicants declared only a single preference despite there being no limits placed on the length of the preference list. This may be a legacy of the previous system in which applicants did not choose a portfolio of schools and in which it was implied

¹⁰This is likely due to several factors: i) students prefer schools closer to home, and schools in more crowded areas were filled under the previous distance-based algorithm; ii) students dissatisfied with their assigned school can request a transfer; iii) applicants may have strategically reported addresses near preferred schools under the prior location-based assignment system.

¹¹The preferences of applicants with siblings already enrolled at their school of interest are detailed in Panel A of Table 14 in Appendix B. We do not have data on applicants whose siblings are enrolled in schools not listed in their preferences.

that applicants were largely assigned to a school based on distance (walking or driving, as obtained from Google Maps) rather than their preferences.

A complementary explanation for the large number of short application lists is that, as shown in Figure 3 of Appendix A, applicants surveyed after the application period indicated high expectations of being assigned to their top choice. These responses were collected before results were published to avoid bias. When asked why they did not add more programs to their rank-ordered list (ROL), 56% of respondents said they lacked information on nearby alternatives, 33% were confident they would be assigned to their top choice, 6% found it difficult to find more schools, 4% preferred no assignment over adding more options, and 1% preferred a non-public school (outside option).

In any case, the fact that the CCAS was new to families in Manta likely also resulted in them not fully adapting their behavior to the new system and rules, meaning that they may not have taken full advantage of the introduction of parental choice and preference reporting. If this is the case, our findings on the welfare gains obtained with the introduction of the CCAS system are probably downward biased when compared with the longer-term results that will eventually be obtained once families are fully accustomed to the new system.

4 Mechanism Result and Welfare Comparison

Our analysis in this section is based on the fact that applicants' reported preference orderings are an accurate representation of true family preferences. This assertion is supported by the non-strategic nature of the DA mechanism, which was furthermore emphasized in the pilot program's communication strategy. We thus compare the share of applicants who were assigned to one of their preferred options under both alternatives, and then use reported preferences to estimate welfare differences. Our welfare analysis forms part of a broader body of literature that uses structural models to study family preferences over school attributes and school choice policy counterfactuals (e.g., Neilson (2021); Kapor et al. (2020); Abdulkadiroğlu et al. (2017); Idoux (2022)).

To estimate preferences, we adopt the utility model of Abdulkadiroğlu et al. (2017) and apply their Markov Chain Monte Carlo (MCMC) estimation method using Gibbs sampling (Rossi et al., 1996), similar to recent applications in the school choice context (Kapor et al., 2020; Idoux, 2022). Our methodology encounters two key challenges: a limited set of school covariates and the assumption that applicants lack full knowledge of all available schools. To address the first, we estimate a model without covariates, where differences in school appeal are driven by a school-specific effect, unobservable to the econometrician but assumed to be known to families. We also introduce a random coefficient for the distance parameter to capture variations in how families value a school's utility relative to distance. For the second challenge, we work with a geographically constrained set of alternatives, assuming families are aware of all schools within their region, though addressing the broader issue of how families form their consideration sets is beyond the scope of this study.

We compare the distance-centric (DC) algorithm described in subsection 2.1 with the Deferred Acceptance (DA) algorithm, using the Stable Improvement Cycles (SIC) (Erdil & Ergin, 2008) and Top Trading Cycles (TTC) (Abdulkadiroğlu & Sönmez, 2003) algorithms as

benchmarks. The TTC algorithm serves as our welfare benchmark, as it delivers a student-optimal assignment, resulting in higher welfare than SIC, which produces a stable but welfare-constrained assignment due to the stability restriction. TTC also outperforms DA, which ensures stability but not student-optimality. To ensure a consistent welfare comparison, we follow the approach in [Abdulkadiroğlu et al. \(2017\)](#), first running the SIC algorithm over each DA assignment, and then applying TTC to the resulting DA-SIC assignment.¹² We conduct 100 lottery simulations for both the DA and DC algorithms to derive our welfare calculations, as in [Abdulkadiroğlu et al. \(2017\)](#).

With regard to the DC algorithm, one relevant point is that parents could strategically report a different address, using someone else’s electricity bill (CUEN) in order to be placed at a preferred school. To include this possibility in the analysis, we run counterfactual assignments in which a random proportion of the applicants strategically choose an address close to their most preferred program. We use different random proportions as we do not have a good estimate of CUEN misreporting under the previous system.

To compare mechanisms, we first re-run the DA algorithm used in the pilot. We use the same inputs, except that we do not include students with special needs in order to make the assignment comparable to that of the DC algorithm.¹³ In the implemented DA, the reported preference rankings were appended to all non-ranked programs using a linear distance sorting criterion. Applicants received a lower priority in the distance-imputed preferences to maximize assignment to the reported preferences.¹⁴ We define assignments to imputed preference as non-preference assignments to distinguish them from the overall assignment obtained with the DC algorithm.

To replicate Ecuador’s previous system (described in subsection 2.1), we consider all available programs and rank them using linear distance sorting. Students with siblings in the system were assigned (if possible) to their sibling’s school before the main process was initiated. To this end, we create a priority group for these students that only applies at the schools in which their siblings are enrolled. This priority is followed by a priority for groups of siblings applying together, as these groups were processed before individual applicants in the main process. This priority is thus applied to all available programs. Finally, given that applicants were processed sequentially, we run a single tie-breaking lottery to break ties.

4.1 Utility Estimation

To estimate welfare, we first need to estimate the parameters determining the utility that families would receive from an assignment to a particular school. To this end, equation 1 presents our utility model.¹⁵ This approach has been increasingly adopted in the literature ([Abdulka-](#)

¹²In this case, SIC and TTC produce the same assignment, meaning no Pareto efficiency improvements are attainable by relaxing stability constraints.

¹³We eliminate both students and vacancies related to special needs, which account for only 0.23% of applicants. This decision was made because students with special needs had a special assignment round before the regular one.

¹⁴When referring to applicant preferences, we intend reported preferences without the distance-imputed preferences.

¹⁵Here, we basically adapt the model used by [Abdulkadiroğlu et al. \(2017\)](#), estimated through Gibbs sampling ([Rossi et al., 1996](#)). We estimate utility using the Markov-Chain Monte Carlo (MCMC), a Bayesian estimation procedure. We therefore use the same conjugate priors, specifically the Inverse-Wishart distribution. The full utility

dirođlu et al., 2017; Kapor et al., 2020; Idoux, 2022), mainly because of its relative ease of implementation in ranked-ordered data contexts.

$$\begin{aligned}
u_{ij} &= S_{ij}\lambda + \delta_j + (-1 + \gamma_i)d_{ij} + \epsilon_{ij} \\
\delta_j &= \bar{\delta} + \xi_j \\
\bar{\delta} &\equiv 0
\end{aligned} \tag{1}$$

$$\begin{aligned}
\gamma_i &\sim \mathcal{N}(0, \sigma_\gamma) \\
\xi_j &\sim \mathcal{N}(0, \sigma_\xi) \\
\epsilon_{i,j} &\sim \mathcal{N}(0, \sigma_\epsilon)
\end{aligned}$$

Here, S_{ij} is a dummy variable equal to 1 if applicant i has a sibling in school j , d_{ij} is the distance between applicant i and school j , and ξ_j is a school-specific preference that is unobservable to the econometrician but that families do observe when comparing alternatives. To identify the parameters, we need to determine a scale normalization for the utility, which we do by setting $\bar{\delta} = 0$ as in Abdulkadirođlu et al. (2017). We also include random coefficients over distance to school, represented by γ_i , in order to consider the heterogeneity in the relative importance of school attributes and travel distance for different applicants.¹⁶ Given that in this utility specification, units of utility are expressed in distance units (km) – which is a result of imposing a -1 parameter on the average dis-utility of linear distance to school – using a random coefficient on distance is quite similar to using a random coefficient for school attributes as in Abdulkadirođlu et al. (2017), in that doing so ends up affecting the relative importance of distance or school attributes. As explained above, by specifying utility in this way, we change the relative relevance of the distance to school and the school-specific unobservable.

The specification in equation 1 assumes that an applicant’s utility increases when a sibling is already enrolled in a school. This, however, does not take into account that the reason for the sibling being enrolled in that school may be because the family liked the school when the sibling enrolled (or transferred) in the first place. This implies that ϵ_{ij} is not random for such cases, highlighting the bias in the estimation. In Appendix B.1 we therefore present our preference parameter estimates and welfare calculations with no sibling-related considerations. For robustness, we further present our results without including the random coefficient on distance. Overall, the findings and conclusions remain the same across these alternate specifications.

As in Abdulkadirođlu et al. (2017), identification relies on the assumption that families report their preferences truthfully and consider all the alternatives within their geographic unit.

specification, including priors, is provided in Appendix D.

¹⁶Conceptually, in each iteration of the Gibbs sampler, utilities are drawn using the estimated parameters of the previous iteration, using reported preference rankings to restrict possible values. Specifically, assuming that i ’s ranking is of size R (1 being the most preferred alternative and R the least preferred), utilities are drawn iteratively using a truncated normal distribution so that:

$$u_{i,j(r=1)} > u_{i,j(r=2)} > \dots > u_{i,j(r=R)} > u_{i,j(r=\hat{r})}, \forall \hat{r} > R$$

To do this iterative sampling, $u_{i,j(r=1)}$ is drawn from $(u_{i,j(r=2)}, \infty)$ if $R > 1$ and using $u_{i,j(r=2)}$ from the previous iteration, and from $(-\infty, \infty)$ when $R = 1$.

Likewise, the key conditional independence assumption is that

$$(\gamma_i, \epsilon_{ij}) \perp d_{ij} | \zeta_j$$

which, in our case, implies that conditional on the vertical school-specific parameter, unobserved tastes for programs are independent of linear distance to school. The previous system, in which families could borrow or buy an electricity bill near their preferred school instead of actually changing their residence, aligns with this conditional independence assumption.

Table 5 in Appendix B presents the estimates from equation 1, and the potential scale reduction factors (Gelman et al., 1992) to assess mixing and convergence of the Gibbs sampling procedure (values close to one imply convergence). We also present the trace plots of the estimated σ_ϵ in each iteration of the Gibbs sampling in Figure 6 of Appendix A. We discarded the initial 50,000 iterations of the Gibbs sampler as a burn period and used the following 100,000 to compute the mean parameters and standard deviations. The trace plots show that the values of σ_ϵ remained stable. Estimates eliminating random coefficients from equation 1 are presented in Table 10 of Appendix B.1, and estimates of the main specification without siblings are presented in Table 11 of the same appendix section. We observe that estimated parameters are very similar across the three alternative models, consistent with the similarity of the welfare estimates using the different specifications.

As in Abdulkadiroğlu et al. (2017), we estimate utilities conditional on the estimated parameters and, importantly, also conditional on the reported preference rankings:

$$\mathbb{E} [u_{i,j} | r_i, \zeta, \lambda, \sigma_\epsilon, \sigma_\zeta, \Sigma_\gamma, d_i]$$

Here, r_i represents i 's reported preference ranking, and to compute i 's expected utility if assigned to school j we directly average over the iterations of the Gibbs sampler procedure, which allows us to easily condition on reported preference rankings. Estimated average utilities are measured in kilometers (km), which is a feature of using a scale normalization of -1 on the linear distance parameter.

5 Results

In this section, we begin by describing the differences between systems in terms of assignment to preferences and linear distance to home. We then present our welfare comparison using the utility model introduced above.

Notably, 55.5% of applicants (n=2,206) reported preferences aligned with the ranking used in the distance assignment mechanism, emphasizing the importance of proximity to families. More than half chose the nearest school (or the closest with a sibling enrolled) as their first preference. Table 2 compares a single simulation of the DA and DC algorithms. Although the DC alternative assigned fewer applicants to their preferred school overall, the percentages were similar for those prioritizing distance, as shown in rows 1 and 2 of Panel B. This indicates that a coordinated system incorporating preferences does not disadvantage applicants focused on proximity. However, the DA algorithm assigned students to schools an average of 0.29 km farther away than the DC algorithm.

Table 2: Mechanism Comparison - Results

	DA	Distance
Panel A: Applicants assigned in:		
<i>Any preference</i>	3,118 (78.44%)	1,986 (49.96%)
<i>First preference</i>	2,773 (69.76%)	1,686 (42.42%)
<i>Average assignment distance</i>	1.30 km	1.01 km
Panel B: Applicants with the same 1st preference (2,206) assigned in:		
<i>Any preference</i>	1,779 (81.16%)	1,537 (70.12%)
<i>First preference</i>	1,644 (75.00%)	1,472 (67.15%)

Figure 4 in Appendix A shows the assignment to different declared preference rankings for both systems. As we can see, the DA algorithm assigns more students to their first preference than the DC algorithm (70% to 42%) and much fewer students to an alternative outside of their reported preference list (22% vs 50%). Tables 6 to 8 in Appendix B and Figure 5 in Appendix A display these results for the different grades. Greater congestion leads to smaller differences between the two mechanisms in terms of applicants being assigned to their preferred options. However, there are two forces at play. On the one hand, more congestion implies that fewer applicants are assigned to a reported preference when using the DA alternative. On the other hand, under the distance-based alternative, more congestion increases the probability that one applicant who is placed in a closer school displaces another who would have ranked that school at the top of their list (particularly in the cases where the latter applicant’s first preference and closest school coincide).

To evaluate the effect of location misreporting, we compute counterfactual assignments in which a random sample of applicants report the location of their most preferred school as their address instead of their true residence. The exercise simulates cases in which families submit another household’s electricity bill to maximize the likelihood of being assigned to their most preferred school. We compute assignments with misreporting levels of 10%, 30%, 50%, 70%, and 90%. These results are available at table 15 in Appendix B. The results of this exercise show that, as the percentage of applicants who change their location increases, the percentage of applicants assigned to one of their preferences rises as well (from 50% to 59%). Nevertheless, the rates of assignment to a preferred option does not reach the level of the DA algorithm, since applicants who misreport their location can only signal a preference for a single alternative.

Table 3 presents the estimated differences in mean utilities (both in km), as well as in standard deviations with respect to a student-optimal (TTC) assignment benchmark.¹⁷ In Panel A, we can see that differences between the DA and TTC algorithms are small in terms of welfare (less than 80 meters) at the preschool levels, and significantly smaller than the welfare loss under the DC (distance) alternative (0.689 km and 0.430 km on average, respectively). The difference is larger when we consider only applicants assigned to different schools under the

¹⁷SIC and TTC actually have the exact same assignment in all 100 simulations, as explained in Appendix E.

different algorithms, as shown in Panel B.

Table 3: Differences in Welfare: Student-Optimal vs. DC and DA Algorithms

Measure	Preschool 1		Preschool 2		Primary 1	
	DC	DA	DC	DA	DC	DA
Panel A: All simulated applicants						
Δ Mean utility (km)	-0.689	-0.003	-0.430	-0.076	-0.318	-0.306
$\frac{\Delta \text{Mean utility}}{\sigma_{Ut, FB}}$	-0.699	-0.003	-0.158	-0.028	-0.094	-0.090
Panel B: Applicants with different assignments across algorithms						
Δ Mean utility (km)	-1.486	-0.005	-0.666	-0.118	-0.417	-0.402
$\frac{\Delta \text{Mean utility}}{\sigma_{Ut, FB}}$	-1.466	-0.005	-0.263	-0.046	-0.127	-0.121

Δ Mean utility (km) is measured computing $u_{i,j(\mu)} - u_{i,j(TTC)}$, where $j(\mu)$ represents the school to which individual i is assigned under mechanism μ . We then compute average utilities for each algorithm and simulation and finally compute the average for each algorithm across simulations. $\frac{\Delta \text{Mean utility}}{\sigma_{Ut, FB}}$ simply uses the utility variance under the TTC mechanism to scale this difference in each simulation. This is done to facilitate extrapolations to other contexts. This same table is presented in Appendix B.1 for the specification without siblings.

In Primary 1, the difference between the DA and the DC algorithms is smaller because more applicants are assigned to a non-preferred alternative due to increased congestion, as shown in Figure 5 of Appendix A. Moreover, in Table 9 of Appendix B, we can see that the share of applicants assigned to the same non-preferred school in both algorithms increases significantly in Primary 1 (56.7% of applicants assigned to the same school, compared to 2.69% and 21.35% in Preschool 1 and 2 respectively). Furthermore, conditional on having a different assignment in the DC and DA algorithms, the share of applicants who move from a non-preferred to a preferred assignment under the DA algorithm is 40.85% in Primary 1, compared to 77% in Preschool 1 and 53% in Preschool 2 (i.e., the share of applicants with improved outcomes is smaller in later years). Finally, the DC algorithm finds on average schools that are closer to home, which is a feature of not prioritizing reported preferences. Given that utility is on average greater for applicants with a lower home-to-school distance, this leads to a lower average difference in utility between mechanisms.

The distribution of estimated welfare overall and in each grade is presented in Figure 7 of Appendix A. Here, we observe that the phenomenon described in the above paragraph occurs in all grades, with two peaks in utility in each figure: one among applicants assigned to a preferred option and another for those assigned to a non-preferred option that is close to home. The DA and TTC (and SIC) algorithms have very similar distributions. However, the TTC algorithm does improve the assignment relative to the DA algorithm in Primary 1, which is explained by the fact that, with higher congestion, stability constraints imposed by tie-breaking lotteries are more restrictive. By eliminating them, TTC (and SIC) achieve a significant improvement (0.318 km overall over the DA assignment and 0.417 km if restricted to applicants with different assignments), as shown in Table 3.

Figure 1: Welfare Differences Between Algorithms (km)




Figure 1 compares welfare gains across different mechanisms, similar to Figure 5 in [Abdulkadiroğlu et al. \(2017\)](#). Although the magnitudes differ between our study and theirs, the proportions are quite similar, suggesting that coordinated mechanisms incorporating applicant preferences yield comparable results in both cases. The differences in magnitude are due to the distinct contexts of New York and Manta, as well as the focus on secondary schools in [Abdulkadiroğlu et al. \(2017\)](#), where applicants are more willing to travel, versus preschool and early primary in our study, where proximity is more valued.

Improvements over the DA algorithm are context-dependent, as seen in the variation across different grades. The potential for improvement is more significant in congested grades (e.g., post-entry-level grades), where the margin is more relevant. A practical approach would be to first implement a CCAS, then analyze the reported preferences to evaluate the potential of SIC, TTC, or other algorithms. This would allow for a careful consideration of the trade-offs between enhancing Pareto efficiency and sacrificing stability or strategy-proofness.

6 Discussion

Both developed and developing countries are increasingly adopting centralized choice and assignment platforms for schools [Neilson \(2021\)](#); [World Bank \(2024\)](#). While there is evidence of the benefits from incorporating family preferences and coordinating assignment in the context of New York City [Abdulkadiroğlu et al. \(2017\)](#), there is limited evidence from developing countries context. In this paper, we study the equilibrium welfare effects of changing the way school seats were allocated in Manta, Ecuador. Prior to the policy change, the system assigned

students ignoring their preferences and solely based on the linear distance between their declared home and schools. We evaluate the impact of adopting a system that collects ranked ordered lists and uses a deferred acceptance (DA) algorithm that is strategy-proof. We follow the empirical strategy used in [Abdulkadiroğlu et al. \(2017\)](#) for New York City and estimate preferences to quantify the welfare effects of the policy change and compare outcomes under different counterfactuals.

Our main result is that implementing a coordinated mechanism that incorporates applicants' preferences yields substantial welfare benefits. This finding aligns with results from New York City even though our developing country setting is very different in many ways.

Specifically we find that the new allocation based on the DA algorithm increases the percentage of applicants assigned to a preferred school from 49.96% to 78.44% and first-choice assignments from 42.42% to 69.76%, with an average increase of 0.29 km in home-to-school distance. Welfare gains range from 0.683 km to 0.012 km, with greater gaps for applicants assigned to different schools under each mechanism.

Our results show that coordinated school choice systems can increase efficiency and benefit families in developing country context. Given the results are similar to those in developed countries, it is likely that the benefits of these systems are generalizable across different contexts.

References

- Abdulkadiroğlu, A., Agarwal, N., & Pathak, P. A. (2017). The welfare effects of coordinated assignment: Evidence from the new york city high school match. *American Economic Review*, 107(12), 3635–89.
- Abdulkadiroğlu, A. & Sönmez, T. (1998). Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3), 689–701.
- Abdulkadiroğlu, A. & Sönmez, T. (2003). School choice: A mechanism design approach. *American Economic Review*, 93(3), 729–747.
- Arteaga, F., Kapor, A. J., Neilson, C. A., & Zimmerman, S. D. (2021). *Smart Matching Platforms and Heterogeneous Beliefs in Centralized School Choice*. Working Paper 28946, National Bureau of Economic Research.
- Correa, J., Epstein, R., Escobar, J., Rios, I., Bahamondes, B., Bonet, C., Epstein, N., Aramayo, N., Castillo, M., Cristi, A., & et al. (2019). School choice in chile. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC '19* (pp. 325–343). New York, NY, USA: Association for Computing Machinery.
- Erdil, A. & Ergin, H. (2008). What's the matter with tie-breaking? improving efficiency in school choice. *American Economic Review*, 98(3), 669–89.
- Gale, D. & Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1), 9–15.
- Gelman, A., Rubin, D. B., et al. (1992). Inference from iterative simulation using multiple sequences. *Statistical science*, 7(4), 457–472.
- Idoux, C. (2022). Integrating new york city schools: The role of admission criteria and family preferences. *Job Market Paper, MIT Economics Department*.
- Kapor, A. J., Neilson, C. A., & Zimmerman, S. D. (2020). Heterogeneous beliefs and school choice mechanisms. *American Economic Review*, 110(5), 1274–1315.
- Neilson, C. A. (2021). Targeted vouchers, competition among schools, and the academic achievement of poor students. *R&R Econometrica*.
- Pathak, P. A. (2011). The mechanism design approach to student assignment. *Annual Review of Economics*, 3(1), 513–536.
- Pathak, P. A. (2017). What really matters in designing school choice mechanisms. *Advances in Economics and Econometrics*, 1(12), 176–214.
- Rossi, P. E., McCulloch, R. E., & Allenby, G. M. (1996). The value of purchase history data in target marketing. *Marketing Science*, 15(4), 321–340.
- World Bank (2024). *World Development Report 2024: The Middle-Income Trap*. Washington, DC: World Bank.

A Figures

Figure 2: Distribution of Declared Applicant Priorities and Ranked Ordered List Size

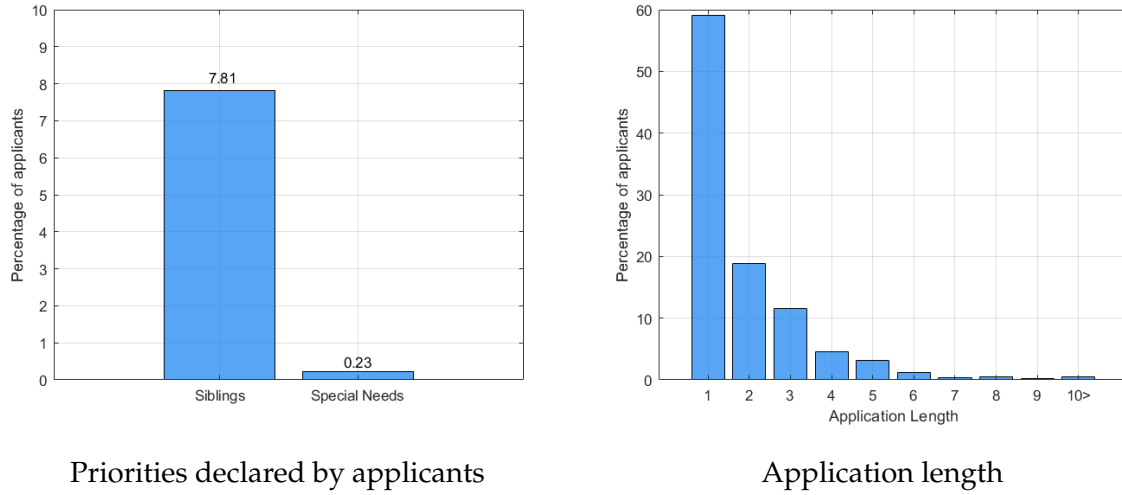



Figure 3: Perceived Probability of Admission to 1st Preference



These responses were obtained in an online survey carried out after the end of the application period but before assignment results were communicated (to avoid biasing responses).

Figure 4: Ranking Assigned: DA and Distance Mechanism






Figure 5: Ranking Assigned by Grade: DA and Distance Mechanism

Preschool 1



Preschool 2



Primary 1






Figure 6: Trace Plots σ_e in Main Specification

Preschool 1



Preschool 2



Primary 1







Figure 7: Welfare Distribution




All applicants



Preschool 1



Preschool 2



Primary 1

In this figure, we plot the utilities obtained with our model when using the scale normalizations $\bar{\delta} \equiv 0$ and -1 as the average disutility from each linear km of distance between the school and the reported location of the family. The level of utility is not relevant, as it depends on the normalization. However, the mass from the utility distribution when using the distance-centric algorithm being shifted to the left is relevant, as it indicates how the relative distributions of utilities compare, and lead to the average differences presented in Table 3.

Figure 8: Figure 5 of Abdulkadiroğlu et al. (2017)

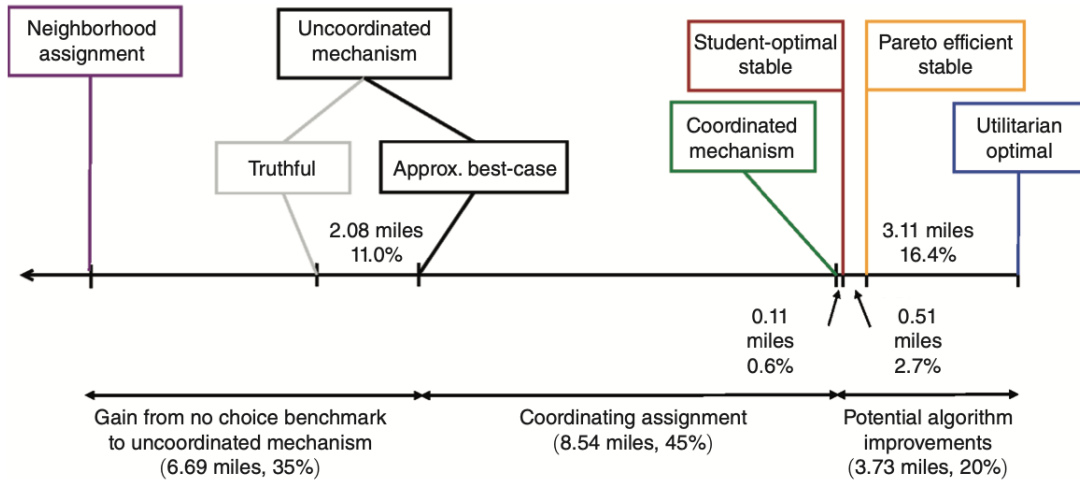


FIGURE 5. COORDINATING ASSIGNMENTS VERSUS ALGORITHM IMPROVEMENTS

B Tables

Table 4: Comparative Statistics for Guayaquil and Manta

	Guayaquil	Manta
Total population	2,291,158	221,122
Population 3-5 years old (% of population 3-17 years old)	19.7	19.1
Minors in the school system (% of population 3-17 years old)	78.9	80.2
Average Mother's Education (of minors 3-17 years old)	11.3 years	10.8 years
Total schools	885	153
Share of public schools	54%	43%
Share of private schools	44%	54%
Share of "fiscomisional" schools	2%	3%
Total enrollment	687,046	86,455
Share of enrollment in public schools	57%	67%
Share of enrollment in private schools	40%	25%
Share of enrollment in "fiscomisional" schools	4%	8%

Table 5: Estimates and Potential Scale Reduction Factors: Main Specification

Estimate	Preschool 1		Preschool 2		Primary 1	
	Mean (SD)	PSRF	Mean (SD)	PSRF	Mean (SD)	PSRF
ξ_j	0.073 (0.492)		-0.918 (0.559)		0.048 (0.648)	
λ	3.055 (0.365)	1.003	4.732 (0.562)	1.007	3.790 (0.842)	1.001
σ_ξ	0.291 (0.066)	1.001	1.251 (0.453)	1.017	0.540 (0.143)	1.004
σ_ϵ	1.042 (0.048)	1.060	1.057 (0.060)	1.025	1.303 (0.102)	1.020
σ_γ	1.253		1.690		1.079	
Tot. schools	55		57		54	
Tot. students	1,098		885		389	

Table 6: Mechanism Comparison - Results Preschool 1

	DA	DC
Panel A: Applicants assigned in:		
<i>Any preference</i>	1,654 (96.95%)	1,102 (64.60%)
<i>First preference</i>	1,592 (93.32%)	966 (56.62%)
<i>Average assignment distance</i>	0.87 km	0.52 km
Panel B: Applicants with the same 1st preference (985) assigned in:		
<i>Any preference</i>	965 (97.97%)	955 (96.95%)
<i>First preference</i>	942 (95.63%)	946 (96.04%)

Table 7: Mechanism Comparison - Results Preschool 2

	DA	DC
Panel A: Applicants assigned in:		
<i>Any preference</i>	1,199 (70.45%)	737 (43.30%)
<i>First preference</i>	93.32 (58.52%)	609 (35.78%)
<i>Average assignment distance</i>	1.32 km	1.08 km
Panel B: Applicants with the same 1st preference (959) assigned in:		
<i>Any preference</i>	708 (73.83%)	607 (63.30%)
<i>First preference</i>	618 (64.44%)	569 (59.33%)

Table 8: Mechanism Comparison - Results Primary 1

	DA	DC
Panel A: Applicants assigned in:		
<i>Any preference</i>	265 (46.74%)	147 (25.93%)
<i>First preference</i>	185 (32.63%)	111 (19.58%)
<i>Average assignment distance</i>	2.56 km	2.29 km
Panel B: Applicants with the same 1st preference (281) assigned in:		
<i>Any preference</i>	144 (51.25%)	112 (39.86%)
<i>First preference</i>	116 (41.28%)	103 (36.65%)

Table 9: Assignment In and Out Preferences under DA and DC Algorithms

	Preschool 1	Preschool 2	Primary 1
Applicants assigned to same schools under DA and DC	967 (56.68%)	801 (47.06%)	261 (46.03%)
Applicants assigned to different schools under DA and DC	739 (43.32%)	901 (52.94%)	306 (53.97%)
<i>Applicants assigned to same schools under DA and DC</i>			
Both DA and DC in preferences	941 (97.31%)	630 (78.65%)	113 (43.30%)
Both DA and DC out of preferences	26 (2.69%)	171 (21.35%)	148 (56.70%)
<i>Applicants assigned to different schools under DA and DC</i>			
DA in preferences and DC out of preferences	569 (77.00%)	478 (53.05%)	125 (40.85%)
DA out of preferences and DC in preferences	17 (2.30%)	16 (1.78%)	7 (2.29%)
Both DA and DC in preferences	144 (19.49%)	91 (10.10%)	27 (8.82%)
Both DA and DC out of preferences	9 (1.22%)	316 (35.07%)	147 (48.04%)

B.1 Appendix Robustness Checks

Table 10: Estimates and Potential Scale Reduction Factors. Main Specification without Random Coefficients

Estimate	Preschool 1		Preschool 2		Primary 1	
	Mean (SD)	PSRF	Mean (SD)	PSRF	Mean (SD)	PSRF
$\tilde{\zeta}_j$	0.131 (0.532)		-1.480 (0.693)		0.075 (0.606)	
λ	3.486 (0.396)	1	4.719 (0.561)	1	4.301 (0.919)	1
$\sigma_{\tilde{\zeta}}$	0.350 (0.082)	1.001	2.766 (0.747)	1	0.497 (0.131)	1
σ_{ϵ}	1.403 (0.054)	1	1.314 (0.061)	1	1.702 (0.121)	1
Tot. schools	55		57		54	
Tot. students	1,098		885		389	

Table 11: Estimates and Potential Scale Reduction Factors. Main Specification without Siblings

Estimate	Preschool 1		Preschool 2		Primary 1	
	Mean (SD)	PSRF	Mean (SD)	PSRF	Mean (SD)	PSRF
$\tilde{\zeta}_j$	0.070 (0.500)		-1.121 (0.559)		0.044 (0.643)	
$\sigma_{\tilde{\zeta}}$	0.300 (0.068)	1.005	1.661 (0.534)	1	0.536 (0.142)	1
σ_{ϵ}	1.089 (0.062)	1.025	0.968 (0.054)	1.038	1.306 (0.101)	1.001
σ_{γ}	1.462		1.376		0.896	
Tot. schools	55		57		54	
Tot. students	1,021		839		345	

Table 12: Differences in Welfare: Student-Optimal vs. DC and DA algorithms. Specification without Random Coefficients

Measure	Preschool 1		Preschool 2		Primary 1	
	Dist	DA	Dist	DA	Dist	DA
Panel A: All simulated applicants						
Δ Mean utility (km)	-0.773	-0.003	-0.456	-0.071	-0.317	-0.296
$\frac{\Delta \text{Mean utility}}{\sigma_{Ut, FB}}$	-0.750	-0.003	-0.177	-0.027	-0.091	-0.085
Panel B: Applicants with different assignments across algorithms						
Δ Mean utility (km)	-1.667	-0.006	-0.707	-0.109	-0.415	-0.388
$\frac{\Delta \text{Mean utility}}{\sigma_{Ut, FB}}$	-1.631	-0.006	-0.302	-0.047	-0.124	-0.116

Δ Mean utility (km) is measured computing $u_{i,j(\mu)} - u_{i,j(TTC)}$, where $j(\mu)$ represents the school to which individual i is assigned under mechanism μ . We then compute average utilities for each algorithm and simulation and finally compute the average for each algorithm across simulations. $\frac{\Delta \text{Mean utility}}{\sigma_{Ut, FB}}$ simply uses the utility variance under the TTC mechanism to scale this difference in each simulation. This is done to facilitate extrapolations to other contexts.

Table 13: Differences in welfare: Student-optimal vs DC and DA algorithms. Specification without Siblings

Measure	Preschool 1		Preschool 2		Primary 1	
	Dist	DA	Dist	DA	Dist	DA
Panel A: All simulated applicants						
Δ Mean utility (km)	-0.717	-0.004	-0.348	-0.083	-0.176	-0.361
$\frac{\Delta \text{Mean utility}}{\sigma_{Ut, FB}}$	-0.756	-0.004	-0.154	-0.037	-0.052	-0.104
Panel B: Applicants with different assignments across algorithms						
Δ Mean utility (km)	-1.560	-0.009	-0.543	-0.130	-0.228	-0.470
$\frac{\Delta \text{Mean utility}}{\sigma_{Ut, FB}}$	-1.442	-0.008	-0.218	-0.052	-0.066	-0.133

Δ Mean utility (km) is measured computing $u_{i,j(\mu)} - u_{i,j(TTC)}$, where $j(\mu)$ represents the school to which individual i is assigned under mechanism μ . We then compute average utilities for each algorithm and simulation and finally compute the average for each algorithm across simulations. $\frac{\Delta \text{Mean utility}}{\sigma_{Ut, FB}}$ simply uses the utility variance under the TTC mechanism to scale this difference in each simulation. This is done to facilitate extrapolations to other contexts.

Table 14: Priorities and Assignments in DA: Potential Improvements for SIC and TTC

	Preschool 1	Preschool 2	Primary 1
Panel A:	<i>Ranking of schools where an applicant has sibling priority(*)</i>		
1st preference	100	157	42
2nd preference	7	3	1
3rd preference	1	2	0
Panel B:	<i>Ranking of DA assignments for applicants with sibling priority below 1st preference</i>		
1st preference	8	3	0
2nd preference	0	2(**)	1(***)
3rd preference	0	0	0

Note: None of the potential applicants that could participate in an improvement cycle (Panel B) coincide in the programs to which they were applying, such that no cycles were attainable.

(*) Panel A shows the highest ranked program where applicants have a sibling priority. If an applicant has priority in both the 1st and 2nd preference, they will only appear in the 1st preference in this table.

(**) One of these two applicants had sibling priority in their second preference, and the other had sibling priority in their third preference.

(***) This applicant had sibling priority in their second preference.

Table 15: Mechanism Comparison with Location Misreporting

	<i>Applicants assigned to any preference</i>	<i>Average Distance</i>
Distance Mech without misreporting	1,986 (49.96%)	1.01 km
Distance Mech + 10% misreporting	2,055 (51.70%)	1.04 km
Distance Mech + 30% misreporting	2,142 (53.89%)	1.09 km
Distance Mech + 50% misreporting	2,230 (56.10%)	1.18 km
Distance Mech + 70% misreporting	2,332 (58.67%)	1.21 km
Distance Mech + 90% misreporting	2,403 (60.45%)	1.29 km
DA Algorithm	3,118 (78.44%)	1.30 km

C Phases of the Distance-Centric Algorithm Implementation Process

The overall process started with the Preparation Phase, in which the Ministry of Education updated all school supply information (i.e., location, available spaces, closure or opening of educational programs, etc.).

In the second, or Registration Phase, families registered their children on a website in order to be granted a spot in a public school. Legal guardians needed to indicate the type of registration (individual or sibling group), the grade level to be attended, any older siblings already enrolled in the public school system, special educational needs, and nationality. They also provided their electricity bill number so as to be geolocated.

This was followed by the Assignment Phase and then the Consultation Phase, during which time families could enter the website to see their school assignments. Finally, the fifth and sixth phases consisted of the School Change Petitions Phase and Continuous Enrollment. Applicants could ask to change schools if there were spaces available, and they could also enroll in a given school once the academic year had already started.

D Full Utility Specification

$$\begin{aligned}u_{ij} &= S_{ij}\lambda + \delta_j - d_{ij} + \gamma_i d_{ij} + \epsilon_{ij} \\ \delta_j &= \bar{\delta} + \xi_j \\ \bar{\delta} &\equiv 0\end{aligned}$$

$$\begin{aligned}\lambda &\sim \mathcal{N}(0, \sigma_\lambda) \\ \gamma_i &\sim \mathcal{N}(0, \sigma_\gamma) \\ \xi_j &\sim \mathcal{N}(0, \sigma_\xi) \\ \epsilon_{i,j} &\sim \mathcal{N}(0, \sigma_\epsilon) \\ \sigma_\gamma &\sim IW(\tau_\gamma, df_\gamma) \\ \sigma_\xi &\sim IW(\tau_\xi, df_\xi) \\ \sigma_\epsilon &\sim IW(\tau_\epsilon, df_\epsilon)\end{aligned}$$

We follow [Rossi et al. \(1996\)](#) and [Abdulkadiroğlu et al. \(2017\)](#) in using disperse priors. The only exception is the use of a smaller τ_γ , given that in this context it is reasonable to impose a smaller prior on the mean variance of the parameter, considering that $\gamma_i > 1$ would imply that a family actually prefers schools farther away from home. Specifically, we use $\sigma_\lambda = 100$, $\tau_\gamma = 0 + size(\gamma_i) = 1$, $df_\gamma = 3 + size(\gamma_i) = 4$,¹⁸ $\tau_\xi = 1$, $df_\xi = 2$, $\tau_\epsilon = 3 + n_{schools}$, and $df_\epsilon = 3 + n_{schools}$.

¹⁸This implies that the mean of the σ_γ prior is 0.5.

E DA-SIC and TTC Equivalence in our Context

As shown in Table 14, there is no potential for priority trading cycles.

The Top Trading Cycles (TTC) algorithm includes the possibility of trading priorities between applicants, which happens when they prefer the alternatives in which they do not have the priority more than ones in which they do, and are thus “willing to trade” the priority. In other words, TTC has the potential to provide improvements over SIC, when there is not a complete correlation between priorities and preferences. In our case, for the priority at declared preferences (over non-preferences imputed by distance), the correlation is one since these are always ranked higher. Thus, the only possibility for the TTC algorithm to improve over the SIC algorithm is to find trades involving the static sibling priority. However, as shown in Table 14 (and explained in the footnote), that is not feasible.

To illustrate this, imagine a system with two schools (A and B), both with only one vacancy, and three applicants (i , j and k). i has priority in A but prefers B over A. j has priority in school B, but prefers A over B. k has priority in both schools, prefers A over B, and has the worst lottery number of the system. The result of the DA and SIC assignment would be i assigned to A and j assigned to B. The TTC algorithm would allow them to trade their priorities and switch their assignments. With that assignment switch, applicant k is now unassigned but has a higher priority in both schools that rejected him (higher priority pre-trade, of course). Such a situation can only arise when the correlation between preference and priority is not one, thus leaving room to trade the priority and get a better assignment.