Designing School Choice for Diversity in the San Francisco Unified School District

M. ALLMAN, I. ASHLAGI, I. LO, J. LOVE, K. MENTZER, H. O'CONNELL, & L. RUIZ-SETZ.

Preliminary draft, please do not distribute.

More than 65 years after the "Brown v. Board of Education" ruling that school segregation is unconstitutional, public schools across the U.S. are resegregating. In attempts to disentangle school segregation from neighborhood segregation, many cities have adopted policies for city-wide choice. However, these policies have largely not improved patterns of segregation. From 2018-2020, we worked with the San Francisco Unified School District (SFUSD) to design a new policy for student assignment system that meets the district's goals of diversity, predictability, and proximity. To develop potential policies, we used optimization techniques to augment and operationalize the district's proposal of restricting choice to zones. We compared these to district-wide choice approaches typically suggested by the school choice literature. We find that appropriately-designed zones with minority reserves can achieve all the district's goals, at the expense of choice, and choice can resegregate diverse zones. Using predictive choice models developed using historical choice data, we show that a zone-based policy can decrease the percentage of racial minorities in high-poverty schools from 29% to 11%, decrease the average travel distance from 1.39 miles to 1.29 miles, and improve predictability, but reduce the percentage of students assigned to one of their top 3 programs from 80% to 59%. Traditional district-wide choice approaches can improve diversity and choice at the expense of proximity. Our work informed the design and approval of a zone-based policy for use starting the 2024-25 school year.

1 INTRODUCTION

One of the barriers to equality of educational opportunity in the U.S. is the re-segregation of schools. Despite the landmark Brown vs Board ruling in 1954 that school segregation is unconstitutional, the concentration of students of color in schools with majority under-represented racial groups has increased nationwide since the 1980s [Orfield and Jarvie, 2020], and school zones across the country have been gerrymandered in ways that perpetuate such segregation [Monarrez, 2018].

The increasing pattern of school segregation along both racial and socio-economic lines is particularly concerning because the largest racial educational achievement gaps often occur when students from racial minorities are concentrated in high-poverty schools [Reardon et al., 2019]. Such achievement gaps are exacerbated by inequitable access to other educational resources, such as private school options, local property taxes or PTA funds for public schools, and experienced teachers.

Motivated by these inequities, we joined with the San Francisco Unified School District (SFUSD) to design and implement a policy for student assignment that enables more diverse enrollment and increased access for students from historically disadvantaged communities. Our work culminated in a policy was approved in December 2020 by the SFUSD Board of Education, and will be implemented in the district starting in the 2024-25 school year.

Traditionally, many school districts use a neighborhood assignment policy, where neighborhood zones are drawn and students attend their neighborhood school. However, neighborhood assignment can reflect local residential segregation in school composition. In the early 2000s, school districts in cities across the US started adopting district-wide choice policies, which allow students to express preferences over schools across the city, and then assign students to schools using a computerized algorithm [see, e.g., Abdulkadiroğlu et al., 2005a,b]. While such choice systems provided families with increased options and allowed some families to attend schools outside their residential area, they also enabled resegregation along neighborhood boundaries or in high-performing schools.

Our approach finds middle ground between neighborhood assignment and district-wide choice. Building off policy concepts developed by SFUSD, we explore policies that split the district into geographic 'zones' and restrict students to choose from schools in their zone. We developed an optimization tool for generating zones that are large enough to overcome residential segregation, but small enough to discourage resegregation through choice. Using historical choice data, we fitted a discrete choice model for how families would rank schools given different geographical zones. We developed a simulation engine that predicts how policies perform on various metrics, which we used to evaluate zone-based policies and compare them to approaches typically suggested by the school choice literature.

We find that geographical zones can be designed to improve the diversity and representativeness of incoming cohorts at each school. However, choice can resegregate diverse zones, and geographical zones can improve on diversity over the status quo only by also using reserves. Together, zones and reserves decrease the percentage of racial minorities in mid-to-high-poverty schools from 29% to 11%, decrease the average travel distance from 1.39 miles to 1.29 miles, and improve predictability, but also significantly restrict choice, reducing the percentage of students assigned to one of their top 3 programs from 80% to 59%. Alternatively, taking the traditional school choice approach, the district's diversity targets can be achieved by retaining district-wide choice and substantially redesigning the priority system. Such policies maintain high levels of choice but significantly increase average travel distance (from 1.39 miles to 1.91 miles) and do not improve predictability.

Given these trade-offs and the SFUSD Board of Education's goals of diversity, predictability, proximity, a zone-based policy for student assignment was proposed for use in SFUSD. We are working with district staff and community stakeholders to implement this policy.

1.1 Related Literature

Our work contributes to the design of school choice systems, which use computerized assignment algorithms, guided by priority rules, to assign students to schools. There is a rich and mature literature on developing and characterizing assignment algorithms [see, e.g., Abdulkadiroğlu and Sönmez, 2003, Balinski and Sönmez, 1999, Kesten, 2010, Ma, 1994, Morrill, 2013]. We contribute to a small but growing literature designing inputs such as priorities and choice restrictions, and quantitatively analyzing their impact on distributional outcomes [Abdulkadiroğlu et al., 2009, Ashlagi and Shi, 2014, Azevedo and Leshno, 2016, Correa et al., 2019, Dur et al., 2018, Escobar and Huerta, 2021, Leshno and Lo, 2021, Shi, 2015]. This paper contributes to the design of choice restriction through zones, and motivated by the San Francisco context raises numerous economic questions and computational challenges for future work.

We augment and operationalize district-proposed policy concepts by drawing on a number of existing tools in the literature on school choice. Algorithms for district-wide choice use priorities at each school to guide the assignment of students to schools [Abdulkadiroğlu and Sönmez, 2003, Shi, 2021]; for example, the widely used Deferred Acceptance algorithm guarantees students with higher priority at a school admission to that school over students with lower priority. We augment the zone policy concept with priorities that provide increased choice to students from historically disadvantaged neighborhoods and underserved minorities. Choice algorithms typically achieve distributional goals such as diversity by reserving seats at each school for different types of students. We use such reserves to guard against resegregation [see, e.g., Abdulkadiroğlu and Grigoryan, 2021, for a survey]. There is also a growing literature on how to optimally restrict students' choices to achieve a district-wide objective such as reducing transportation costs [Ashlagi and Shi, 2014, 2016, Dur et al., 2018, Shi, 2021]. We similarly restrict student choice to carefully-selected zones in order to improve diversity.

Our work also builds on the literature designing zones for neighborhood schools [Caro et al., 2004, Clarke and Surkis, 1968, Franklin and Koenigsberg, 1973, Liggett, 1973, Shirabe, 2009]. To our knowledge, we are the first to provide an approach for integrating zone optimization with choice.

2 CASE STUDY: THE SAN FRANCISCO UNIFIED SCHOOL DISTRICT

The San Francisco Unified School District (SFUSD) is one of many districts experiencing school resegregation [Orfield and Jarvie, 2020]. SFUSD serves more than 56,000 students in the San Francisco metro area, and each year assigns around 5,000 kindergarten students to 159 programs at 72 different elementary schools. Students are currently assigned to general education (GE) programs, language programs, and special education programs via district-wide choice.

2.1 Current Assignment Policy

Since 2010, the elementary level at SFUSD has used district-wide choice, guided by 58 contiguous 'attendance areas' that serve as the neighborhood for one school.[1] In January, families submit an unrestricted ranked-order-list (ROL) with their preferences over programs for which they are eligible.[2] In March, the ROLs, program capacities, and a set of priorities are fed into an algorithm,

[1] The remaining 14 schools are city-wide and have no corresponding neighborhood.

[2] Language and special education programs have eligibility criteria, including native or non-native speakers of English, or students with special needs.

which determines the assignment of students to programs. Students who are not assigned to a school on the ROL are 'designated' to their closest unfilled school.[3]

Priorities and lotteries are used to determine who is assigned to over-demanded schools. Highest priority is given to students with a sibling at the program, followed by students living in the 20% of census tracts with the lowest test scores in the 2010 census (CTIP1),[4] and finally students from the school's attendance area. A random lottery, drawn independently at each program, breaks remaining ties, in a process known in the literature as multiple tie-breaking.

Between the years 2010-2018, the assignment algorithm was based on a modification of the Top Trading Cycles algorithm, which allows families to trade priorities in trading cycles for admission to preferred schools [Abdulkadiroğlu and Sönmez, 2003].[5] We noted that the implemented algorithm was neither strategyproof nor envy-free,[6] and families reported confusion regarding the algorithm and its properties. Since 2019, in line with our recommendation, the algorithm has been based on the Deferred Acceptance algorithm [Abdulkadiroğlu and Sönmez, 2003, Gale and Shapley, 1962], which is strategyproof and terminates in an envy-free outcome.

2.2 Challenges

The existing system provided families a wide range of choices, but replicated patterns of residential segregation at SFUSD schools. In U.S. public schools, eligibility for free or reduced-price lunch (FRL) is often used as a measure of socioeconomic need. While 52% of SFUSD elementary school students are eligible for FRL, more than 70% of schools had a percentage of students eligible for FRL that differed from the district average by more than 15%. Figure 1 shows that schools with concentrations of students with higher socio-economic need are located in neighborhoods with higher socio-economic need. A similar pattern of neighborhood and school segregation occurs with African American, Latinx and Pacific Islander (AALPI) students.

The existing choice system also created complications for families. We note that families vary widely in the number of programs on their ROLs. In 2018-2019, White, Asian, Latinx and African American students listed on average 15.9, 7.3, 6.1 and 3.4 programs, respectively, with median list lengths of 10, 4, 3, and 2 programs. Similarly, there are substantial differences based on socio-economic status: students from neighborhoods with greater than 50% FRL eligibility rank 5.7 schools on average, while students from neighborhoods with less than 50% FRL eligibility rank 14.2 on average, with median list lengths of 3 and 9 respectively.[7] These differences hold similarly across groups based both on race and socio-economic status, controlling for CTIP1 eligibility.[8] Shorter lists limit a family's opportunities to be assigned to a school of their choice. In addition, families expressed concerns about the gameability and transparency of the system, and uncertainty and anxiety due to the large number of choices.

---

[3]A second, similar, round of assignment is conducted in June, with approximately 1,200 families. This includes roughly 300 families who did not participate in the first round, and families seeking to change their assignment from the first round. There are a few more rounds of assignments for families based on requests for a single school or special needs.

[4]In a highly gentrifying area, CTIP1 may result in assigning priority to unintended families.

[5]The description of the algorithm used in SFUSD from 2010-2018 was vague; cycles were chosen to maximize the number of improvements for students.

[6]A school assignment algorithm is strategyproof if families cannot improve their assignment by misreporting their preferences, and envy-free if a student with lower priority at a program is never assigned over a high-priority student who desires to attend the program over their assigned program.

[7]CTIP1 students rank 4.5 schools on average, while non-CTIP1 students rank 11.1, with medians of 3 and 6 respectively. While CTIP1 students tend to live in neighborhoods with high FRL eligibility, they rank fewer programs as their district-wide priority allows the vast majority of CTIP1 students to be assigned to one of their top choices.

[8]Anecdotal evidence suggests that such differences stem from differing information about schools, time and resources to engage with student assignment, understanding of the process, and needs for long lists (e.g. due to CTIP1 priority).

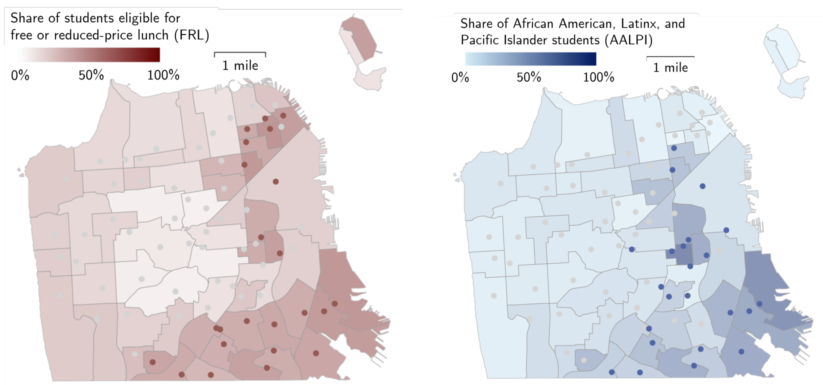

(a) Socio-economic segretation: High FRL

(b) Racial segregation: High AALPI

Fig. 1. Neighborhoods and schools with (a) high socio-economic need, as measured by the percentage of students who are eligible for free or reduced price lunch (FRL), and (b) a high percentage of African American, Latinx, and Pacific Islander (AALPI) students. Each shaded piece represents an attendance area for a school. In Figure 1a, the red dots represent the 50% of schools with higher percentage FRL, and the grey dots represent the 50% of schools with lower percentage FRL. In Figure 1b, the blue dots represent the 50% of schools with higher percentage AALPI, and the grey dots represent the 50% of schools with lower percentage AALPI.

San Francisco also has many private and charter school options, with nearly 20% of families who participate in public school choice ultimately opting out in favor of private or charter school options [9] This is one the highest percentages of any city nationwide.

2.3 Towards a New Assignment Policy

In the fall of 2018, the San Francisco Board of Education passed a resolution "Developing a Community Based Student Assignment", initiating a redesign of the elementary school student assignment process. The redesign process had three primary goals: diversity, proximity, and predictability, using a lens of equity of access.

The *diversity* goal aims to find a student assignment system that creates schools that are socio-economically and racially representative of the district as a whole. A specific focus is ensuring that racial minorities are not concentrated in high-poverty schools [Reardon et al., 2019]. The *proximity* goal aims to give every student the opportunity to go to a school close to home. This reduces travel time, alleviates transportation concerns and costs, and improves community cohesion. The *predictability* goal aims to improve the student assignment experience for families, mitigating families' frustrations with the high uncertainty in district-wide choice, and the fact that their child could be assigned to any school in the district. These objectives require trade-offs; for example, the proximity and diversity goals are often in opposition due to residential segregation.

The District proposed 3 initial policies to begin the redesign process and to solicit community feedback. The concepts were centered around restricting choice to zones.

Concept 1. Students are automatically assigned to their neighborhood school. Families may then opt in to an optional choice process to access all other programs in the district.

Concept 2. Families are restricted to choose programs at schools in a small zone around their home. Zones consist of 3-5 schools and their surrounding neighborhoods. Families also have the option to rank citywide programs, such as language or K-8 programs.

[9]About 7% are assigned to their top choice program but eventually do not enroll.

Concept 3. Families are restricted to choose programs at schools in a medium-sized zone consisting of 8-12 schools. No choice is allowed beyond the zone; all citywide programs are split between zones to ensure families can access all the types of programs SFUSD offers.

3 THEORY TO POLICY

In December 2019, we started working with SFUSD staff to propose a new policy for elementary school student assignment that better met district goals.

3.1 Policy Development

The literature on school choice suggests that priorities, quotas and reserves, and restricting choice can be used by policymakers to shape the distributional outcomes of schools. We augmented the concepts proposed by the concept by building on existing theory through zone design, and bringing the insights from existing literature to the policy discussion.

3.1.1 Restricting Choice via Zones. The district-proposed concepts center around improving predictability and proximity by restricting choice. Existing school choice literature has used optimal priority design to inform how best to restrict choice [Shi, 2021]. The Boston Public Schools system was designed using such an approach, giving students personalized menus based on their residential address [Shi, 2015]. SFUSD decision-makers advocated for restricting choice using zones rather than a home-based plan, as they believed zones were more straightforward to communicate to families and easier to update in the event of further residential demographic shifts. Furthermore, zones could potentially align transportation needs, allowing for more efficient use of the district's busing budget. The Board of Education also communicated the necessity for contiguous, visually compact zones, even if such zones restricted the ability to achieve proximity or diversity. The focus on compact, contiguous zones added significant computational complexity to the policy design process, limiting our ability to exactly optimize policy decisions.

To create zones containing multiple schools, we built on optimization techniques previously used to decide on geographic areas for neighborhood schools. These techniques piece together geographic neighborhoods to form zones; we used the 58 contiguous 'attendance areas' as our neighborhoods. To create a set of M zones, let Z be a set of M labels, with each $z \in Z$ corresponding to a zone. Let $x_{z,j}$ be a binary decision variable indicating whether neighborhood $j \in J$ is assigned to zone z. Further, allow c_j to represent the total school capacity within neighborhood j and n_j to represent the total number of students in neighborhood j. The objective is to create contiguous zones satisfying zone size, school capacity, and student diversity constraints, and which best ensure each zone can accommodate all of its students by minimizing the shortage of seats in each zone.

Let B be the maximum allowed difference between the number of students in each zone. The shortage objective and balance constraints are given by

$$
\begin{aligned}
\max \quad & y \\
\text{s.t.} \quad & \sum_j (c_j - n_j) \cdot x_{j,z} \geq y \quad \forall\, z, \\
& \sum_{w \in Z, w \neq z} \left(\sum_j n_j x_{j,z} - \sum_j n_j x_{j,w} \right) \leq B \quad \forall\, z, \\
& \sum_{w \in Z, w \neq z} \left(\sum_j n_j x_{j,z} - \sum_j n_j x_{j,w} \right) \geq -B \quad \forall\, z.
\end{aligned}
$$

For a linear diversity measure, let a_j be the value of the diversity measure within neighborhood j and A be the average value across the district. Let δ_a and Δ_a be (multiplicative) upper and lower bounds for each measure as a proportion of the district average. We considered diversity metrics including percentage of students eligible for FRL, an aggregate score for neighborhood socioeconomic status, and proportion of students belonging to historically disadvantaged racial groups. Diversity constraints are given by

$$\sum_j a_j n_j x_{j,z} \geq \delta_a A \cdot \sum_j n_j x_{j,z} \quad \forall\, z,$$
$$\sum_j a_j n_j x_{j,z} \leq \Delta_a A \cdot \sum_j n_j x_{j,z} \quad \forall\, z.$$

A similar approach was used to create a balance of school characteristics across zones, measured using historical Level 1 test scores (lowest achievement level) and an aggregate score from the California School Dashboard that includes current test performance, chronic absenteeism, and suspension rates as well as trends in these measures.

A zone is *contiguous* if the geographic area defined by the zone is connected. Zone contiguity is desirable for keeping zones simple and interpretable, and also for supporting within-zone transportation. To ensure contiguity of the constructed zones, we built on the approach of Shirabe 2009 [Shirabe, 2009], which pre-specifies zone centers and requires each neighborhood in a zone to be connected to the center via other neighborhoods in the zone. We implemented several heuristics to select centers, including random selection, using the schools with highest test scores or other measures of school quality. Specifically, suppose each zone corresponds to a pre-specified neighborhood, $Z \subseteq J$. Let $N_{j,z} = \{j' \mid d_{j',z} < d_{j,z}\}$ be the set of neighborhoods j' that are adjacent to j and closer to the zone z than j. Then zone contiguity is enforced by

$$x_{j,z} \leq \sum_{k \in N_{j,z}} x_{k,z} \forall j \neq z.$$

For each pair of neighborhoods j, j' let $d_{j,j'}$ be the distance between the two neighborhoods. The following constraint ensures that the average distance between a student's attendance area and the zone center (as defined in the contiguity constraint) is at most D,

$$\sum_j d_{j,z} n_j x_{j,z} \leq D \cdot \sum_j n_j x_{j,z} \quad \forall\, z.$$

Finally, to ensure that the solution corresponds to a feasible assignment with each area corresponding to a single zone, we include the following constraints:

$$\sum_z x_{j,z} = 1 \;\forall\, j, \text{ and } x_{j,z} \in \{0, 1\} \quad \forall\, j, z.$$

To develop a zone policy, we generated many zones via the optimization formulation by varying constraint parameters. Types of constraint variations are described in Appendix B. Next, we simulated the assignment process, restricting choice to within the zones, using the policy evaluation framework outlined in Section 4. We identified zones on the Pareto frontier of the diversity and proximity measures after the choice process (specific metrics are outlined in Section 4.3). Leveraging their greater contextual knowledge, the district selected maps from this smaller, well-performing set. We illustrate the impact of a zone policy with the zone map in Figure 2, which performs well on diversity measures. However, the proposed policy does not specify the zone map: the implemented zone map will be selected through further optimization and community engagement.

Through this zone selection process, we found that small zones (4-6 neighborhood schools) generally are unable to overcome residential segregation patterns. Larger zones (9-11 schools) tended to experience greater choice-driven re-segregation and longer travel distances. Medium-sized zones (6-8 schools), which divide the city into roughly 9 sections, best balance between residential segregation driven and choice driven segregation effects. We provide more details in Appendix C. Based on these findings, to evaluate zone-based policies we use the contiguous 9-zone map in Figure 2.

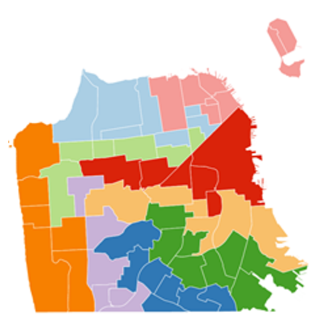

Fig. 2. One set of nine contiguous zones selected for evaluating potential policies. These zones were selected due to performing well on diversity metrics and appearing 'reasonable' to district decision-makers.

3.1.2 Priorities. Priority systems are widely adopted to implement social goals such as proximity and diverse outcomes. A growing number of papers in the school choice literature study how districts can better achieve social outcomes by designing appropriate priorities [Ashlagi and Shi, 2014, 2016, Dur et al., 2018, Shi, 2021]. As SFUSD had already leveraged priorities in the current assignment system, we suggested exploring the ability of priorities to achieve district goals.

The implementation of priorities faces practical challenges. Legal restrictions prohibit the use of individual student characteristics to be used in student assignment.[10] Consequently, standard practice is to use a geographic proxy based on neighborhood characteristics, including socio-economic status, racial composition, and historical academic performance. In our implementation of priorities, we adhered to such practice by assigning priority to students based on the historical characteristics of the city block in which they resided. We focused on neighborhood information that is already available to the District and could reasonably be updated on a regular basis. Although the priorities used in policy implementation are yet to be finalized, in our preliminary analysis we used an index including multi-year averages of lowest-level standardized test score rates, percentage of students from historically underserved ethnic groups, free and reduced price meal eligibility rates, and neighborhood socioeconomic status (including median household income, adult educational attainment, and poverty rate). Details on the index's construction can be found in Appendix A.

The number of different priority groups also impacts student outcomes; we found that 3 priority groups allowed the group with lowest priority to retain sufficient choice while still achieving improvements in diversity. Based on our simulations, we limited our policy space to assignments where priority is given by grouping students into 3 equal-sized groups.

3.1.3 Quotas and Reserves. The school choice literature typically uses quota and reserve systems to achieve distributional goals such as diversity (see, e.g. [Biró et al., 2010, Echenique and Yenmez, 2014, 2017, Ehlers et al., 2014, Erdil and Kumano, 2013, Fragiadakis and Troyan, 2017, Hafalir et al.,

[10]While eligibility for free or reduced lunch (FRL) is possible to use, the district avoided using it in the assignment due to time required to verify it for all students (FRL is used for prioritization in Wake County, NC).

2011, Kamada and Kojima, 2015, Kominers and Sönmez, 2016]). Based on existing literature and precedence at other school districts such as Berkeley Unified, we suggested considering a policy that included reserves to achieve diversity goals.

Similar to the priority structure, we develop a geographic proxy to capture student characteristics and focus on easily accessible data. Again, no specific reserve structure has been finalized by the district, but for this work we use 3 equally-sized groups using the same diversity index used for priorities, described in detail in Appendix A.

3.2 Selected Policies

We selected three policies to evaluate that represent the trade-offs between different policy tools.

Zones. The Zones policy aligns the most closely with the District's proposed Concepts 2 and 3. Students are restricted to choose from within one of nine contiguous zones shown in Figure 2, based on their home location. In addition, motivated by the school choice literature, the Zones policies uses priorities. Students are grouped into 3 groups and prioritized for schools in their zone according to their group.

Zones + Reserves. The Zones + Reserves policy closely mimics the Zones policy, also using the zones in Figure 2. In addition, drawing on the literature on diversity in school choice, the policy adds a reserve system to improve diversity. Based on their neighborhood, students are assigned to one of three diversity categories. Seats are reserved for students of each category proportional to their prevalence in that zone.

Priorities. The Priorities groups students into 3 groups and prioritizes students at all schools according to their group. This policy is the most similar to the preexisting policy and typical district-wide choice approaches, as it gives families access to all programs and attempts to achieve distributional goals using priorities.

We note that the selected policies do not capture the notion of automatic assignment suggested by SFUSD's Concept 1. Initial simulations showed that automatic assignment to a neighborhood school did not significantly change the properties of the assignment relative to current practice, particularly showing no improvement in diversity, so we eliminated this policy from further consideration.

4 POLICY EVALUATION

To evaluate the impact of potential policies, we developed an end-to-end simulation engine to estimate what student assignments might look like under a given policy. The simulation engine took policy decisions as inputs, used the zone optimization formulation to create zones, used historical choice data to predict student choices, simulated the assignment of students to schools, then computed metrics evaluating how well the resulting assignment aligned with district goals.

4.1 Choice Modeling

We used a choice model to predict how students would rank programs under counterfactual policy scenarios. We trained a rank-ordered multinomial-logit (MNL) model using historical students' rank lists from 2014-2018, as well as student and program features present in the data provided by the district. [11]

Let θ represent a student type, and let $X_{\theta,p}$ represent characteristics of the student-school pair. The student's utility for the program p is given by

$$u(\theta, p) = \beta_p^\top X_{\theta,p}.$$

[11]Since the assignment policy during these years was not strategyproof, we also estimated choices based on data from 2019-2020, where the policy switched to a strategyproof assignment mechanism; the estimates were similar.

Then, the probability that a student of type θ chooses program p as their most preferred program (denoted Y_θ) is given by

$$\mathbb{P}[Y_\theta = p] = \frac{e^{u(\theta,p)}}{\sum_{p' \in P} e^{u(\theta,p')}}.$$

We identified a set of predictive features by evaluating the predictive power of school and program features previously used in the school choice modelling literature. For greater flexibility, we considered some non-linear functions of these features as well. From these options, we only retained features with statistically significant predictive power to avoid over-fitting. The features we used for our model were the square root of the distance from the student to the program, whether the program is within walking distance from the student (less than 0.5 miles), whether the program is in the attendance area of the student, whether the student's home language matches the language of the program if the program is language immersion, the program type. We also included a fixed effect for the school containing the program to capture the multiplicity of factors that can affect the desirability of a school. Furthermore, we separated the square root distance feature for students with and without CTIP1 status because the CTIP1 status students' choices were found to be more sensitive to distances. Feature weights and standard errors are shown in Table 1.

Table 1. Parameter Values for the Choice Model

Feature	Parameter
Sqrt dist CTIP	-3.0
	(0.12)
Sqrt dist non-CTIP	-2.8
	(0.06)
Dist $\leq 0.5 miles$	0.035
	(0.105)
Sibling	6.44
	(0.14)
In att. area	0.93
	(0.06)
Lang. match	1.28
	(0.09)

To validate our choice model, we compare how predicted first choices from the choice model compare to observed first choices from the historical data. We focus on three components of modeling accuracy: distance, program type, and market share. For distance, we show the distribution of distances from students' home location to their first choice program for both historical and modeled preferences. For program type, we observe the frequency of each type of program (general education, Spanish immersion, Cantonese immersion, etc.) in first choices across historical and modeled preferences. Finally, we measure the market share of each program, as measured by the fraction of applications that rank the specific program first. The results of these analyses are shown in Figure 3.

The choice model generates a list over all eligible programs for each student. In practice, however, students tend to submit much shorter ranked lists. Anecdotally, families find it costly to research schools and formulate preferences, leading to abbreviated preference lists. To create preference rankings more representative of historical data, we truncate preference model rankings to a more realistic length. We tested several different protocols for setting list lengths, including constant list

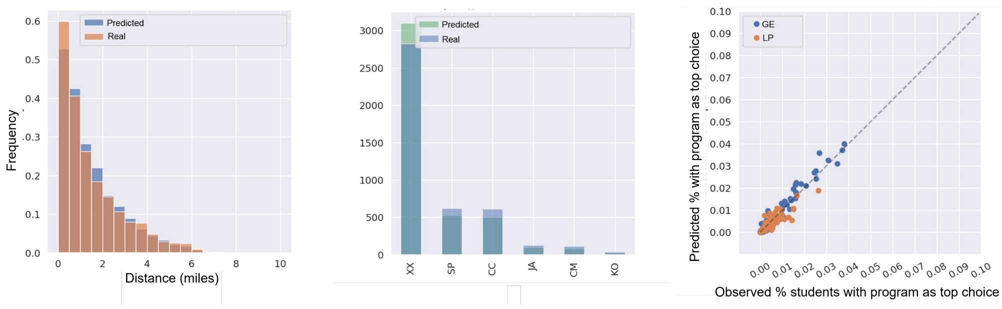

(a) Distance to first choice (b) Program type of first choice (c) Market share of first choice

Fig. 3. Comparisons of students' predicted first choices and their real choices. For plot (b), XX denotes non-language programs, SP denotes Spanish programs, CC denotes Cantonese, JA denotes Japanese, CM denotes Mandarin, and KO denotes Korean. For plot (c), LP represents a language program, and GE represents a general education program.

lengths, lengths determined by ethnicity averages, and list lengths matching the lengths submitted in actual preferences. We evaluate each procedure by comparing the outcome of the 2018-19 student assignment process using our choice model and list length scheme against the real outcome in the historical data. We found that ranking seven programs led to assignments closest to the real outcome.[12]

4.2 Assignment Mechanism

Students are assigned to programs using the student-proposing Deferred Acceptance (DA) algorithm with designation (Algorithm 1). DA takes as input students' preferences $\succ^{\theta}$, program capacities q_p, as well as program priorities $\succ_p$. all of which can be affected by the assignment policy. The algorithm generates an assignment of students to schools that is *envy-free*, meaning that if a student prefers a different school to her own assignment, then she has a lower priority to that school than all students that are currently assigned there. Reserves are implemented via a standard modification of the mechanism [Hafalir et al., 2013].

SFUSD uses an additional step from the standard algorithm to reduce the number of students who are unassigned after the algorithm completes. During this phase known as 'designation', students who have exhausted their preference list are assigned to the nearest school that still has empty seats. Students are considered in a random order generated by a lottery number.

ALGORITHM 1: Student-Proposing Deferred Acceptance with Designation

Data: Student preferences $\succ^{\theta}$, program priorities $\succ_p$, program capacities q_p
Result: An assignment of students to schools
while *some student can still propose* **do**
  1. All students propose to their most-preferred program for which the student is eligible that has not previously rejected them;
  2. All programs tentatively accepts the highest priority students (up to capacity) that have proposed to them and reject the rest;

end
Each student is assigned to the program that still tentatively accepts them, or remains unassigned.
Designation: Unassigned students are randomly ordered and sequentially assigned to the nearest program not yet at capacity for which they are eligible.

[12]See [Pathak and Shi, 2021] for similar patterns modeling choice in Boston Public Schools.

4.3 Metrics

To evaluate the resulting student assignments, we selected and developed a number of measures of proximity, predictability, and diversity of the assignment of students to schools. We worked closely with experts on educational inequality and district staff to select metrics that best reflected the priorities of district stakeholders.

Proximity metrics include average distance to assigned school, and percentages of students assigned within a 'walking distance' of 0.5 miles, or assigned further than 3 miles from home.

Diversity measures focused on both socio-economic and racial diversity. One measure captures the fraction of high-poverty schools: the percentage of schools where the FRL average of students assigned to the school is greater than 10% or 15% above the fraction of students eligible for FRL district-wide. A second metric aims to capture the concentration of historically underserved ethnic groups in high poverty schools, measuring the percentage of AALPI students who are assigned to +10% or +15% FRL eligibility schools (AALPI in +10% or +15% FRL schools). Another notion of diversity is representativeness, and we measure this using the percentage of students from different socioeconomic or racial backgrounds that would need a different assignment for every school to be representative of the district (dissimilarity). We also measure isolation of racial minorities, using the percentage of GE programs that have 1-4 African American students (GE with 1-4 AA).[13]

Predictability is measured by the average percentage of programs for each student where admission came down to a random lottery (unpredictable).

Choice measures include the percentage of students who are unassigned or designated, and the number of students receiving one of their top 1 or 3 choices. Note that these top 1 or 3 are defined over all schools, not just programs they are eligible for under the policy (i.e., any of the top 3 choices may not be in their zone). We also measure the percentage of students assigned at least 3 miles away from their home to a program they ranked 5th or higher; this captures students assigned to a school far away that is not one of their top choices. For the 2018-19 assignment, we include students who were designated as not receiving a top choice, since some students ranked fewer than 5 options.

Community cohesion is measured using the percentage of students assigned to the same school as at least 3 students from their neighborhood (defined by a census block group, BG Cohesion (3)).

We measure equity of access by comparing choice, proximity, and diversity metrics for students of different races (African American, Asian, White, Latinx) and socio-economic status (from a neighborhood with high FRL eligibility, i.e. more than 50% of students eligible, or from a neighborhood with low FRL eligibility, i.e. fewer than 50% eligible).[14] We calculate percentage of each subgroup receiving one of their top 3 choices, average distance to assigned school, and the percentage of each subgroup attending a moderate-to-high poverty school (FRL eligibility greater than 15% above the district average). We also calculate the percentage of each subgroup assigned at least 3 miles away from their home to a program they ranked 5th or higher.

4.3.1 *Visualization Tool.*

To further compare policies, we developed a visualization tool to show the outcomes and trade-offs illuminated by the simulation process. The tool was intended for internal use, including district staff and policy decision-makers. We show example figures in Appendix E to demonstrate the functionality of the tool.

[13]In 2018-19, 30.7% of kindergarten applicants identified as African American, Latinx, or Pacific Islander. Due to the relatively high percentage of these students, the district was most concerned about creating schools with a high concentration of students from these historically underserved ethnic groups. African American students made up 4.8% of the applicant class. The relatively small number of these students leads to greater concern about isolation.

[14]In addition, we calculated all metrics for Pacific Islander applicants, but due to the small number of students (18 in the 2018-19 kindergarten applicant class) these values are not as meaningful and we do not report them.

5 RESULTS

We simulated assignment policies using both historical preferences and the estimated choice model. We sampled 100 sets of lottery numbers and simulated each policy with the same lottery numbers but different realizations of the noise in the choice estimation. In Table 2 we report all measures using the estimated choice models and averaged over the 100 iterations.

5.1 Data

We report simulations based on applications to kindergarten for the 2018-19 school year.[15] In this school year there were 5153 student applying for kindergarten placement in 159 programs in 72 schools in SFUSD, including 65 general education programs, 64 language programs and 30 special education programs. To better capture true demand for SFUSD seats, we removed 381 students who received their first choice program in SFUSD yet did not enroll the following academic year. All simulations were performed using the remaining 4772 students. Due to legal restrictions on student-level socioeconomic data, we are only provided with average eligibility for FRL over 4 years for each census block. To examine student sub-groups, we categorize any student from a block with more than 50% of students eligible as 'High FRL' and less than 50% eligible as 'Low FRL.'

5.2 Policy Comparison

We compare the policies along the district goals of diversity, proximity, and predictability, and district lens of equity of access. In line with the school choice literature, we also compare the performance of the policies on choice and community cohesion.

5.2.1 Proximity. Across all measures of proximity, Priorities performs significantly worse than the other policies. For example, the average distance of a student's assigned program from their home location is significantly farther for Priorities at 1.91 miles, while the other policies are under 1.4 miles. Zones and Zones + Reserves provide a moderate improvement over 2018-2019 in the average distance traveled and the fraction of students traveling over 3 miles. However, the 2018-2019 assignment allows the largest fraction of students to travel under 0.5 miles. This suggests that the distribution of travel distances in the 2018-2019 assignment is more dispersed than for Zones and Zones + Reserves.

5.2.2 Diversity. Across most measures of diversity, Priorities and Zones + Reserves perform the best, and significantly improve upon 2018-2019, which performs the worst. For example, compared to the 2018-2019 assignment, Zones + Reserves and Priorities approximately halve the percentage of AALPI students attending schools where the fraction of students eligible for FRL is greater than 15% above the district average, decreasing it from 20% to 11% and 9% respectively. In general, Priorities and Zones + Reserves perform relatively similarly along diversity metrics, showing that similar levels of diversity are achievable through both the design of priorities as well as through restricting choice. In the metrics where they differ, Zones + Reserves performs significantly better: it achieves lower dissimilarity along socio-economic status compared to all the policies, and Priorities results in a higher number of schools with isolated African American students compared to all the policies.

5.2.3 Predictability. Zones and Zones + Reserves result in many fewer unpredictable GE programs compared to Priorities. This is not surprising, as both policies restrict the GE programs that students can rank. Zones + Reserves leads to a greater number of unpredictable GE programs than Zones, because the reserves ensure students from each Diversity Category are able to be assigned to each program.

[15] We choose the 2018-19 school year as it was the most recent year wherein SFUSD used the student assignment system that had been in place since 2011.

		Zones	Zones + Reserves	Priorities	2018-2019
Proximity	*Avg Distance (miles)*	1.35	1.29	1.91	1.39
	Distance ≤ 0.5 miles	30%	32%	24%	34%
	Distance ≥ 3 miles	10%	10%	23%	14%
Diversity	*+10% FRL Schools*	27%	20%	23%	31%
	+15% FRL Schools	15%	11%	9%	20%
	AALPI in +10% FRL	40%	25%	31%	42%
	AALPI in +15% FRL	24%	11%	11%	29%
	Dissimilarity AALPI	39%	35%	33%	37%
	Dissimilarity SES 3	35%	23%	30%	30%
	Schools with 1-4 AA	64%	66%	71%	66%
Predictability	*Unpredictable GE*	5%	9%	40%	
Choice	*Unassigned*	1.8%	1.5%	0.2%	1.9%
	Designated	11%	11%	19%	12%
	Rank Top 1	37%	34%	60%	61%
	Rank Top 3	64%	59%	74%	80%
	Dist >= 3, Rank >= 5	5%	6%	10%	4%
Community Cohesion	*BG Cohesion (3)*	10%	9%	6%	8%

Table 2. Average assignment metrics using preferences generated by choice model for counterfactual policies, and the real 2018-2019 student match.

5.2.4 Choice. Priorities and the 2018-2019 assignment give students significantly higher access to their top-ranked choices than Zones and Zones + Reserves, with Zones giving slightly higher choice than Zones + Reserves. For example, Zones and Zones + Reserves reduces the percentage of students assigned to one of their top 3 programs from 80% in 2018-2019 to 64% and 59% respectively. Unsurprisingly, restricting students' choices within zones reduces their chances of being assigned to one of their top-ranked programs, which may lay outside their zone. Reserves further reduce choice on average, but the effect is smaller than the introduction of zones. Priorities allow students to retain access to their top-ranked choices, but also increases the percentage of students assigned low in their list. Specifically, the Priorities policy significantly increases designation from 12% to 19%,[16] and also significantly increases the percentage of students assigned to a school more than 3 miles away and ranked 5th or worse, from 4% to 10%.

5.2.5 Community Cohesion. The community cohesion metric follows a similar pattern to the average distance metric, with Priorities having the lowest cohesion and 2018-2019 the second lowest. The differences across policies are small.

5.2.6 Equity of Access. Table 3 reports assignment metrics by student racial groups and by FRL eligibility. On the whole, Zones and Zones + Reserves lead to worse choice metrics but better proximity metrics for all subgroups compared to the 2018-19 assignment and the Priorities policy.

[16]This increase is notable even if we include both designated and unassigned students: compared to 13.9% in 2018-2019, this number slightly decreases to 12.8% and 12.6% with Zones and Zones + Reserves respectively, and significantly increases to 19.2% with Priorities.

Table 3. Equity of Access Measures describing choice, proximity, and concentration in higher poverty schools by student group.

		Zones	Zones + Reserves	Priorities	2018-19
African American	*Rank Top 3*	67%	57%	92%	94%
	Avg Distance (miles)	1.45	1.27	2.24	1.82
	In + 15% FRL	16%	20%	9.1%	26%
	Dist >=3, Rank >= 5	3.4%	3.5%	2.6%	1.2%
Asian	*Rank Top 3*	70%	65%	78%	83%
	Avg Distance (miles)	1.43	1.43	1.97	1.56
	In + 15% FRL	15%	16%	9.8%	26%
	Dist >= 3, Rank >= 5	6.8%	8.3%	11.1%	3.7%
Latinx	*Rank Top 3*	75%	62%	89%	90%
	Avg Distance (miles)	1.26	1.26	1.71	1.34
	In + 15% FRL	26%	8.9%	11%	29%
	Dist >= 3, Rank >= 5	2.1%	4.1%	3.3%	1.2%
White	*Rank Top 3*	50%	55%	56%	85%
	Avg Distance (miles)	1.33	1.18	1.87	1.09
	In + 15% FRL	5.6%	3.1%	3.5%	5.7%
	Dist >= 3, Rank >= 5	6.6%	4.9%	15.1%	2.2%
Low FRL	*Rank Top 3*	48%	54%	51%	76%
	Avg Distance (miles)	1.38	1.25	2.01	1.26
	In + 15% FRL	4.0%	2.9%	2.6%	7.0%
	Dist >= 3, Rank >= 5	8.5%	6.6%	19.4%	5.1%
High FRL	*Rank Top 3*	78%	64%	96%	85%
	Avg Distance (miles)	1.32	1.33	1.83	1.52
	In + 15% FRL	24%	16%	13%	36%
	Dist >= 3, Rank >= 5	2.3%	5.0%	1.5%	2.7%

We note that African American and Latinx students experience a greater drop in choice than White or Asian students. From the 2018-19 assignment to the Zones + Reserves policy, the Rank Top 3 metric drops by 37% for African American students, 31% for White students, 28% for Latinx students, and 18% for Asian students. All subgroups except White students experience an improvement in distance metrics from the 2018-19 assignment to the Zones + Reserves policy. Comparing the Priorities policy to the current policy, White students experience the largest reduction in choice metrics, likely driven by an increase in the number of students receiving an equity priority. The Priorities policy offers the highest choice metrics across the 3 evaluated policies for all subgroups, but also has the highest average distance to assigned schools.

Although both Zones + Reserves and Priorities achieve improved diversity measures in aggregate, we see that the effect of these policies varies by subgroup. Noting the percentage of each subgroup that is assigned to a moderate-to-high poverty school (+15% FRL above district average), we see that adding reserves significantly decreases the proportion for Latinx students, but increases slightly the fraction of African American students. Much of the gain from adding reserves in the aggregate metric, the percentage of AALPI students assigned to +15% schools, comes from the improvement of the Latinx student group from 16% to 8.9%. Since there are significantly more Latinx students than African American students, this obscures a slight worsening from 16% to 20% for African American students. By contrast, the Priorities policy assigns far fewer African American students to moderate-to-high poverty schools, only 9.1% compared to Zones + Reserves' 20%.

The final metric, *Dist >= 3, Rank >= 5*, captures the fraction of students who are doing poorly on both proximity and choice. For focal groups (African American, Latinx, and High FRL), the Priorities policy reduces the percentage students who are traveling far to a school that is not one of their top choices relative to the 2018-19 assignment, but this percentage increases for all other subgroups. By contrast, Zones + Reserves shows a uniform but relatively small increase for all subgroups relative to the 2018-2019 assignment. The difference between these policies is most pronounced for non-focal families; the Zones + Reserves policy has 6.6% of students in low-FRL neighborhoods traveling far to a less preferred school, whereas the Priorities policy has 19.4%.

In addition to sub-group analysis, we developed a rough benchmark of the proximity versus choice trade-offs that different student sub-groups face, conditioned on (nearly) representative diversity. Such a benchmark helps capture differences across groups that lie outside student assignment policy, such as the geographic distribution of schools, differing geographic distribution of students from different groups, and differences in student preferences over schools. For example, certain historically underserved groups, like African American or Pacific Islander students, have a much wider spatial distribution of top choices compared to other student groups, potentially due to perceived higher quality schools being located in areas of the city that are further away from these communities. The benchmark was implemented using a minimum weight matching formulation, and is described in detail in Appendix D.

The benchmark provides an approximate Pareto frontier for how well our selected policies are performing on distance and choice. All three policies achieve distance and choice metric values that are close to different regions of the Pareto frontier. The Priorities policy is closer to the choice-maximal regime, while Zones and Zones + Reserves are both closer to the proximity maximal regime. Although further work is required, preliminary analysis suggests that while no policy performs well on both choice metrics and proximity metrics, this is a necessary trade-off if we require diversity. Choosing a location on the Pareto frontier represents a value decision on the part of the policymaker.

5.3 Policy Decision

Taken together, our results reveal clear trade-offs between diversity, proximity, and choice. While Zones + Reserves and Priorities both improve diversity compared to the 2018-19 assignment, the former does so at the expense of choice, while the latter does so at the expense of proximity.

Moreover, these trade-offs differ in notable ways across subgroups of students. While, compared to the 2018-19 assignment, both Zones + Reserves and Priorities significantly reduce the percentage of African American, Latinx, and Pacific Islander students assigned to mid-to-high poverty schools, under the Zones + Reserves policy this is driven primarily by a reduction for Latinx students, and African American students are still concentrated to a similar extent in mid-to-high poverty schools. If the goal of diversity is to reduce the concentration of racial minorities in high-poverty schools, this suggests that Priorities may be preferable to Zones + Reserves.

From the perspective of choice, while the Priorities policy allows more students to be assigned to their top-ranked programs than the Zones + Reserves policy, compared to all other policies it also significantly increases the percentage of students assigned to a school far away that is not one of their top choices. At the subgroup level, this is true for White and Asian students, and students living in blocks with fewer than 50% of students eligible for FRL, while the comparison between Priorities and Zones + Reserves is reversed for African American, Latinx, and High FRL families. If one goal of the policy change is to avoid parental dissatisfaction and reduced enrollment in SFUSD schools, this suggests that Zones + Reserves may be preferable to the Priorities policy. On the other hand, the Priorities policy may be preferable for reducing undesirable assignments for African American and Latinx students, and students with socio-economic need.

Together with the SFUSD Ad Hoc Committee on Student Assignment, we presented these trade-offs to the SFUSD Board of Education. The Board determined that its main goals were diversity, proximity, and predictability, and that diversity was the most important district goal. Notably, while the district is committed to a lens of equity of access, choice is not a district goal. As a result, the approved policy guidelines are based on the Zones + Reserves policy. The approved policy uses medium, contiguous zones to restrict choice. Each zone must have a student population that falls within 15% of the district average for FRL eligibility. Reserves, termed 'Guard Rails' by the district, will further be used to improve school representatives of district demographics, and will be defined using district-defined diversity categories. An additional priority referred to as an 'equity tiebreaker' will be determined to allow target student populations (e.g. homeless students, students in foster care) better access to their desired schools.[17]

The SFUSD Board of Education adopted the policy in December 8, 2020. We will continue to work with district staff and stakeholders to determine the exact zone boundaries, incorporating multiple rounds of community feedback, further optimization, transportation analysis, and input from demographers into the design process. We are also working with district staff to specify guard rails and eligibility for the equity tiebreaker.

6 CONCLUSION AND DISCUSSION

In this work, we contribute to the school choice literature by exploring the intersection of zones and choice. While zone design has been a longstanding component of school desegregation [Clarke and Surkis, 1968, Liggett, 1973], previous attempts typically do not account for individual preferences.[18] This paper illustrates an approach to combine zones with choice to impact student assignment distributional outcomes and student experiences. Our findings illustrate that considering both underlying population characteristics and school choice patterns is necessary to achieve diverse school assignments and mitigate choice-driven re-segregation. We provide a method for incorporating choice in zone optimization, but note that our simulations are not necessarily generating an "optimal policy" due to heuristic approximations.

Our second contribution is to integrate optimization and simulation tools to suggest and evaluate potential policies and support policy-making decisions. In close partnership with district collaborators, we developed a end-to-end policy-support tool that takes policy settings as inputs, simulates assignments, and evaluates and visualizes performance on district goals. This tool was central to policymakers' decisions regarding student assignment trade-offs and ultimately led to development of the policy recommendation.

Limitations. The selection of a zone-based policy was influenced by the district goals of predictability, proximity, and diversity over choice. The policy search space evaluated in this work is not exhaustive; rather, it reflects the preferences of district decision-makers for zone-based policies over alternative methods of restricting choice.

The restriction to a zone-based policy also creates challenges with assigning students close to home. Under the zone-based policy, around two thirds of all families will be assigned to a school further than a walk-zone of half a mile from their home, an increase of approximately 2% from

[17] The equity priority is not intended to further the district's diversity goals.

[18] See, e.g., the assignment policy in Berkeley https://www.berkeleyschools.net/wp-content/uploads/2020/06/Elementary-Student-Assignment-Administrative-Regulation.pdf.

the 2018-19 status quo. Families near the ends of zones may need to travel further than in other choice-restricting policies, such as the home-based policy in Boston. Many students may also find the shift to zones takes away access to schools that are near their place of residence. District support for the policy relies on the premise that zones can reduce transportation costs and improve the district's ability to provide busing for students. However, SFUSD currently provides limited general education busing, and largely relies on public transportation and private vehicles.[19] A true improvement in proximity necessitates greater SFUSD investment in transportation.

Determining priorities and reserves for the policy suffer from similar drawbacks as CTIP1. These are based on block-level characteristics rather than individual ones. This is especially challenging in San Francisco, which exhibits high levels of gentrification.[20] Some districts, like Wake County in NC, use individual eligibility for FRL to determine priority. Operational and legal challenges prevent such individual characteristics from being used in the SFUSD assignment process.

Restricting Choice. A clear drawback of zone-based policies are the limits they place on choice. While choice was not an explicit re-design goal, there are several reasons a policy that severely limits choice may be of concern. First, an overarching district objective is to ensure all families have access to a school meeting their unique needs, and a primary way for families to express their needs is through choice. Second, limiting choice may harm the communities of color the district deems focal.[21] Community engagement revealed that focal families were concerned about policies that reduced their choice, and our simulations suggest these communities will see the highest reduction in choice. Shifts in programs and resources will be required to ensure focal families receive the policy as serving their needs.

In addition, around 20% of families who participate in the SFUSD assignment system eventually do not enroll, and enrollment in SFUSD public schools has further decreased during 2020-2022 due to the global pandemic.[22] If fewer families receive what they perceive to be a desirable school, the district may experience further enrollment decreases. Diversity can be harmed if families from certain groups, such as those of higher socio-economic status, increase the rate at which they leave the system. Our study does not account for students opting out in the counterfactual policies, and policies that negatively affect choice may enhance this negative trend in enrollment and perform worse than suggested by our simulations. As enrollment at public schools heavily influences future funding, decreased enrollment due to choice dissatisfaction could negatively impact the ability of SFUSD to resource educational opportunities in the medium to long term.

One alternative to choice that has been taken in the literature, and which was used in student assignment in Boston Public Schools (BPS), is ensuring students have access to a range of schools according to certain quality measures [Shi, 2015]. However, while in BPS quality was central to the district policy goals, SFUSD preferred to avoid the use of schools' quality measures in the design and evaluation of zones.[23]

The paradigm proposed by the district is that changing a student assignment policy can lead to changes in choice patterns and perceived school quality in the medium- to long-term. The district believes that a zone plan will allow greater community investment in neighborhood schools

[19]SFUSD busing is guaranteed only for the approximately 2.7% of students who qualify for transportation through special education eligibility.

[20]In Berkeley, which uses a similar approach, there is a clear pattern for dividing the cities into zones and determining diversity categories.

[21]The district's focal families are African American, Latinx, Pacific Islander, and low-income families.

[22]Note that 7% receive their top choice and do not enroll. This is likely due to timing of the assignment decisions in the different education systems.

[23]This preference is due to the fact that the designation of quality schools may incentivize families to herd at such schools, thus proving a self-fulfilling prophecy.

through greater community cohesion, eventually leading to the cultivation of desired programs closer to home for all students. In this work, in line with the scope of the literature on school choice and mechanism design, we take student choice patterns to be fixed, and do not incorporate such longer-term behavioral shifts. Community building and strategic investment will be necessary to ensure that these potential zone benefits are realized.

Next Steps. The proposed policy institutes significant changes to the assignment process, and the success of the policy relies on community buy-in. Using reserves to improve diversity and restricting choice using zones are both new in the SFUSD context. It will thus be crucial to extensively engage with and communicate decisions and trade-offs clearly to families, and incorporate community feedback into the implementation of the student assignment policy. The district is currently engaging families to help increase understanding of the proposed policy and collect feedback. We are working to support final implementation decisions, including selecting zone boundaries, which will be decided jointly with families and other stakeholders. We are also working on community engagement and informational tools that can help reduce disparities in families' information about schools, school assignment, and involvement in implementation decisions.

A location-based policy can also prompt residential sorting and gentrification. The policy implementation will need to minimize the harms of such residential sorting, and be flexible to account for demographic shifts over time. In the policy decision, SFUSD committed to regularly reevaluating the system, allowing for the adjustment of zone boundaries and reserves. We will also work with the district to empirically evaluate the policy's impact on residential sorting and gentrification.

The problem of disentangling residential and school segregation exists in many districts beyond the SFUSD context. Beyond zones, there is a wide range of student assignment policies that admit more computationally tractable optimization formulations [see, e.g. Shi, 2021], and which may better be able to achieve distributional goals in other districts. We leave the design of such policies to future work.

Finally, diversity and representation are not equivalent to integration. School choice alone cannot fully address educational disparities that arise from neighborhood and school segregation. Decision-makers at SFUSD are hopeful that the new zone-based policy will enable community investment in all zone schools, rather than letting a small number schools be perceived to be most desirable or of highest quality. Community engagement that builds trust in the system will also enable more families, teachers, and other stakeholders to see their role in enabling integration. We are also engaging with the district on providing cost-effective transportation routes within zones, and using resource allocation techniques to redistribute resources such as programs and facilities more equitably across the district. Ultimately, our work with the district on diversifying schools is just the first step in enabling more equitable access to education.

REFERENCES

Atila Abdulkadiroğlu and Aram Grigoryan. 2021. School Choice. (2021).

Atila Abdulkadiroğlu, Parag A Pathak, and Alvin E Roth. 2005a. The new york city high school match. *American Economic Review* 95, 2 (2005), 364–367.

Atila Abdulkadiroğlu, Parag A Pathak, and Alvin E Roth. 2009. Strategy-proofness versus efficiency in matching with indifferences: Redesigning the NYC high school match. *American Economic Review* 99, 5 (2009), 1954–78.

Atila Abdulkadiroğlu, Parag A Pathak, Alvin E Roth, and Tayfun Sönmez. 2005b. The Boston public school match. *American Economic Review* 95, 2 (2005), 368–371.

Atila Abdulkadiroğlu and Tayfun Sönmez. 2003. School choice: A mechanism design approach. *American economic review* 93, 3 (2003), 729–747.

Itai Ashlagi and Peng Shi. 2014. Improving community cohesion in school choice via correlated-lottery implementation. *Operations Research* 62, 6 (2014), 1247–1264. https://doi.org/10.1287/opre.2014.1319

Itai Ashlagi and Peng Shi. 2016. Optimal allocation without money: An engineering approach. *Management Science* 62, 4 (2016), 1078–1097.

Eduardo M Azevedo and Jacob D Leshno. 2016. A supply and demand framework for two-sided matching markets. *Journal of Political Economy* 124, 5 (2016), 1235–1268.

Michel Balinski and Tayfun Sönmez. 1999. A tale of two mechanisms: Student placement. *Journal of Economic theory* 84, 1 (1999), 73–94.

Péter Biró, Tamás Fleiner, Robert W Irving, and David F Manlove. 2010. The college admissions problem with lower and common quotas. *Theoretical Computer Science* 411, 34-36 (2010), 3136–3153.

F. Caro, T. Shirabe, M. Guignard, and A. Weintraub. 2004. School redistricting: Eembedding GIS tools with integer programming. *Journal of the Operational Research Society* 55, 8 (2004), 836–849. https://doi.org/10.1057/palgrave.jors.2601729

S Clarke and J Surkis. 1968. An operations research approach to racial desegregation of school systems. *Socio-Economic Planning Sciences* 1, 3 (1968), 259–272.

Jose Correa, Rafael Epstein, Juan Escobar, Ignacio Rios, Bastian Bahamondes, Carlos Bonet, Natalie Epstein, Nicolas Aramayo, Martin Castillo, Andres Cristi, et al. 2019. School choice in Chile. In *Proceedings of the 2019 ACM Conference on Economics and Computation*. 325–343.

Umut Dur, Scott Duke Kominers, Parag A Pathak, and Tayfun Sönmez. 2018. Reserve design: Unintended consequences and the demise of Boston's walk zones. *Journal of Political Economy* 126, 6 (2018), 2457–2479.

Federico Echenique and M Bumin Yenmez. 2014. Online appendix to how to control controlled school choice. (2014).

Federico Echenique and M. Bumin Yenmez. 2017. How to control controlled school choice. *American Economic Review* 107, 4 (2017), 1362–1364. https://doi.org/10.1257/aer.20160913

Lars Ehlers, Isa E Hafalir, M Bumin Yenmez, and Muhammed A Yildirim. 2014. School choice with controlled choice constraints: Hard bounds versus soft bounds. *Journal of Economic Theory* 153 (2014), 648–683.

A. Erdil and T. Kumano. 2013. Prioritizing Diversity in School Choice.

Juan F Escobar and Leonel Huerta. 2021. School choice and segregation. (2021).

Daniel Fragiadakis and Peter Troyan. 2017. Improving matching under hard distributional constraints. *Theoretical Economics* 12, 2 (2017), 863–908.

Allen D Franklin and Ernest Koenigsberg. 1973. Computed school assignments in a large district. *Operations Research* 21, 2 (1973), 413–426.

David Gale and Lloyd S Shapley. 1962. College admissions and the stability of marriage. *The American Mathematical Monthly* 69, 1 (1962), 9–15.

Isa Hafalir, M. Yenmez, and Muhammed Yildirim. 2011. Effective Affirmative Action in School Choice. *Theoretical Economics* 8 (04 2011). https://doi.org/10.2139/ssrn.1806084

Isa E Hafalir, M Bumin Yenmez, and Muhammed A Yildirim. 2013. Effective affirmative action in school choice. *Theoretical Economics* 8, 2 (2013), 325–363.

Yuichiro Kamada and Fuhito Kojima. 2015. Efficient matching under distributional constraints: Theory and applications. *American Economic Review* 105, 1 (2015), 67–99.

Onur Kesten. 2010. School choice with consent. *The Quarterly Journal of Economics* 125, 3 (2010), 1297–1348.

Scott Duke Kominers and Tayfun Sönmez. 2016. Matching with slot-specific priorities: Theory. *Theoretical Economics* 11, 2 (2016), 683–710.

Jacob D Leshno and Irene Lo. 2021. The cutoff structure of top trading cycles in school choice. *The Review of Economic Studies* 88, 4 (2021), 1582–1623.

Robin Segerblom Liggett. 1973. The Application of an Implicit Enumeration Algorithm to the School Desegregation Problem. *Management Science* 20, 2 (1973), 159–168. https://doi.org/10.1287/mnsc.20.2.159

Jinpeng Ma. 1994. Strategy-proofness and the strict core in a market with indivisibilities. *International Journal of Game Theory* 23, 1 (1994), 75–83.

Tomas Monarrez. 2018. Attendance boundary policy and the segregation of public schools in the united states. *Unpublished manuscript. May* 4 (2018).

Thayer Morrill. 2013. An alternative characterization of top trading cycles. *Economic Theory* 54, 1 (2013), 181–197.

Gary Orfield and Danielle Jarvie. 2020. Black Segregation Matters: School Resegregation and Black Educational Opportunity. *Civil Rights Project-Proyecto Derechos Civiles* (2020).

Parag A Pathak and Peng Shi. 2021. How well do structural demand models work? Counterfactual predictions in school choice. *Journal of Econometrics* 222, 1 (2021), 161–195.

Sean F Reardon, Ericka Weathers, Erin Fahle, Heewon Jang, and Demetra Kalogrides. 2019. Is separate still unequal? New evidence on school segregation and racial academic achievement gaps. (2019).

Peng Shi. 2015. Guiding school-choice reform through novel applications of operations research. *Interfaces* 45, 2 (2015), 117–132.

Peng Shi. 2021. Optimal priority-based allocation mechanisms. *Management Science* (2021).

Takeshi Shirabe. 2009. Districting modeling with exact contiguity constraints. *Environment and Planning B: Planning and Design* 36, 6 (2009), 1053–1066. https://doi.org/10.1068/b34104

A DIVERSITY INDEX

We constructed new diversity priorities using a diversity index defined by district staff. The diversity index is defined for each census block group, and is an average of 4 scores: a FRL score, a neighborhood SES score, an academic score, and AALPI score (African American, Latinx and Pacific Islander).

Let B be the set of all census block groups. For a census block group b, the diversity index is given by a combination of four scores

$$\text{DiversityIndex}(b) = \frac{\text{FRLScore}(b) + \text{SESScore}(b) + \text{AcademicScore}(b) + \text{AALPIScore}(b)}{4}. \quad (1)$$

In (1), the FRL score measures the percentage of students in block b eligible for free or reduced price lunch ($FRL\%(b)$) as given by SFUSD Student Nutrition Services, normalized by the maximum percentage over all blocks,

$$\text{FRLScore}(b) = \frac{FRL\%(b)}{\max_{b' \in B} FRL\%(b')}.$$

The neighborhood socioeconomic status score ($SESScore(b)$) uses data from the American Community Survey 5-year estimates 2013-17, including median household income in the block ($HHInc(b)$), poverty level ($Pov\%(b)$), and adult educational attainment ($BachDeg\%(b)$), as measured by the percentage of residents 25 years of age or older in the block who have a bachelors degree,

$$\text{SESScore}(b) = \frac{\left(1 - \frac{HHInc(b)}{\max_{b' \in B} HHInc(b')} + \frac{Pov\%(b)}{\max_{b' \in B} Pov\%(b')} + 1 - \frac{BachDeg\%(b)}{\max_{b' \in B} BachDeg\%(b)}\right)}{\max_{b'' \in B} \left(1 - \frac{HHInc(b'')}{\max_{b' \in B} HHInc(b')} + \frac{Pov\%(b'')}{\max_{b' \in B} Pov\%(b')} + 1 - \frac{BachDeg\%(b'')}{\max_{b' \in B} BachDeg\%(b')}\right)}.$$

The block group academic score ($AcademicScore(b)$) measures the percentage of students with level 1 test scores,[24] normalized by the maximum percentage over all blocks,

$$\text{AcademicScore}(b) = \frac{L1\%(b)}{\max_{b' \in B} L1\%(b')}.$$

The AALPI score measures the percentage of students from the historically underserved ethnic groups of African American, Latino, and Pacific Islander students,

$$\text{AALPIScore}(b) = \frac{AALPI\%(b)}{\max_{b' \in B} AALPI\%(b')}.$$

In reporting results on student subgroups, we also sometimes group students using an SES index $SESIndex(b)$. Our SES index uses only the socioeconomic and free and reduced price lunch components of the diversity index, and is defined as follows,

$$\text{SESIndex}(b) = \frac{\text{FRLScore}(b) + \text{SESScore}(b)}{4}.$$

In the 3 Diversity Categories priorities, each census block group is designated to Diversity Category 1, 2, or 3 based on whether the block falls in the 1-33rd percentile, the 34-66th percentile, or the 67-100th percentile of a diversity index (with a higher diversity index value indicating greater disadvantage). Each student's Diversity Category is based on the census block group of their residence. Students in Category 1 have first priority, followed by students in Category 2 and then Category 3.

[24]There are five different achievement levels, Level 1 is the lowest achievement level.

B OPTIMIZATION PARAMETERS

Throughout the zone design process, we have implemented a variety of constraints that capture different desirable characteristics of a zone policy. At the time of writing, the district has not selected a specific zone map for the policy, and we are developing additional zone optimization tools to better inform the zone policy. The following section describes the types of constraints used and varied within the zone generation phase described in Section 3.1.1.

We implemented a number of variations in student representativeness constraints. We attempted bounding the difference in number of students between zones, as well as bounding the difference between zone averages and the district average for several student characteristics. One such characteristic was the proportion of students in each zone that are African American, Latinx, or Pacific Islander (one of the district's definitions of focal families). Another race-based measure of representativeness was that the percentage of students belonging to one of the larger ethnic groups in the district (White, Asian, or Latinx) could not deviate beyond a 10-15% of the district-wide average. Zones could also not deviate more than 10 or 15% from the district average free and reduced price meals eligibility rate. Finally, we also bounded the deviation of the diversity index described in Appendix A compared to the district average.

School characteristics were also considered in zone optimization. We implemented a constraint to ensure that each zone had approximately the same number of schools within a margin of 1-2 schools. We also implemented several versions of school quality constraints in a manner similar to the student representativeness constraints–bounding the deviation from the district average. We considered both English Language and Math standardized test scores, as well as what fraction of students in the school historically have met statewide educational standards. Another school quality metric we considered is the "Color Score" of schools in the California School Dashboard from the California Department of Education. In later exploration, we focused on this final measure because it incorporates school improvement in its evaluation and is thought to be relatively less reinforcing of problematic factors in evaluating school quality.

We have also implemented several different approaches to limiting the geographic spread of zones. Initial attempts used a pairwise constraint on the distance between blocks assigned to the same zone, but this quadratic number of constraints significantly impacted runtime. Alternative distance limits included the cover distance constraint, which limits the student-weighted average distance to the zone centroid, implemented with both euclidean distance and driving travel time. Borrowing from the gerrymandering literature, we have added penalty term on squared distance into the objective, which encourages compactness. We also added an additional penalty term that dis-incentivizes long and winding zone perimeters.

The optimization parameters used to generate the zone map for the Zones and Zones + Reserves policies are as follows. The number of zones was set to 9, and contiguity was to be enforced. Each zone could have a shortage of no more than 168 seats within the neighborhood schools. No zone could differ in number of students from another zone by more than 341 students. The number of schools in each zone could differ from an even distribution of schools to zones by more than a 2 school increase or a 1 school decrease. The student-weighted average distance to the center of the zone could not exceed 3 miles. The FRL multiplicative constraints included a lower bound of 0.4 and an upper bound of 1.5 (though the resulting zones fell within tighter bounds, with estimated FRL eligibility per zone varying only from 31% to 54%). We also included a lower bound on the diversity index described in Appendix A of 40% of the district average, and a lower bound on the AALPIScore (defined in the same section) of 40% of the district average. Centroid schools were randomly selected from a list of schools with historically high academic performance.

C ZONE SIZE

We used the optimization problem in Section 3.1.1 to generate small, medium, and large zones. Small zones divided the city into 13 regions, typically containing 4-6 schools. Medium zones divided the city into 9 zones, each with 6-11 schools. Large zones divided the city into 6 regions containing 9-15 schools. The cover distance constraint parameters considered included $D \in \{3.3, 3.5, 3.8, 4.2\}$ for small zones, $D \in \{3.7, 4, 4.2, 4.5\}$ for medium zones, and $D \in \{3.9, 4, 4.2, 4.5, 4.7\}$ for large zones. We considered both tight and loose balance constraint parameter values, varying from $B = 180$ to $B = 341$ for medium zones, $B \in \{434, 589\}$ for large zones. Some zones were generated using a feasibility model where shortage was considered a constraint, and these parameters varied from 120 to 228. Both shortage and balance parameters were informed by a fraction of the expected number of students per zone.

By varying constraint parameters and methods of selecting zone centers, we generated contiguous and non-contiguous small, medium, and large zones. After removing duplicate output, we had 1521 large contiguous zones, 24 large non-contiguous zones, 130 medium contiguous zones, 28 medium non-contiguous zones, 40 small contiguous zones, and 73 small non-contiguous zones.

The size of zones impacts the feasible diversity and proximity levels of each zone. We found that optimizing for small zones required less restrictive diversity constraint parameters to find a feasible solution. Solutions for large zone optimization could be found with more restrictive diversity constraints. However, large zones required looser distance constraints compared to small zones to find a feasible solution.

In particular, small zones led to worse pre-choice diversity measures compared to medium and large zones for both socio-economic and ethnic diversity. The lack of diversity persists through the choice process, leading to non-diverse school assignments even compared to the current assignment system. Figure 4 shows that across many small zones, none outperform the current assignment policy on a measure of socio-economic diversity: the percentage of schools where the fraction of assigned students eligible for FRL is within 15% of the district average. We concluded that small zones were unable to achieve the adequate underlying zone diversity to achieve the district's goals.

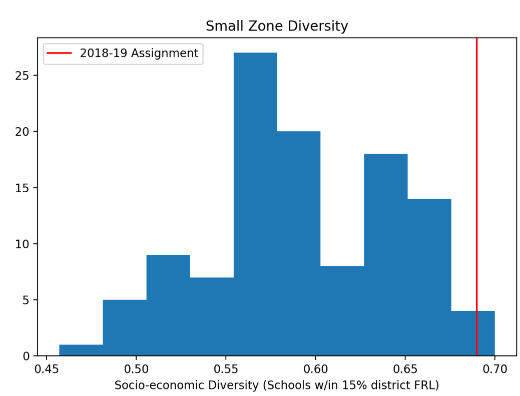

Fig. 4. Across many small zone policies, small zones typically lead to less diverse school assignments than the 2018-2019 assignment.

Although larger zones achieve better diversity, if zones are too large, students have enough choice for some choice-driven re-segregation to occur. Population diversity alone does not guarantee diverse school outcomes after the choice process. As Figure 5a shows, large zones typically outperform smaller zones on pre-choice metrics of racial and socio-economic diversity, but Figure

5b shows that after the choice process much of that diversity advantage disappears when we observe post-choice diversity metrics.

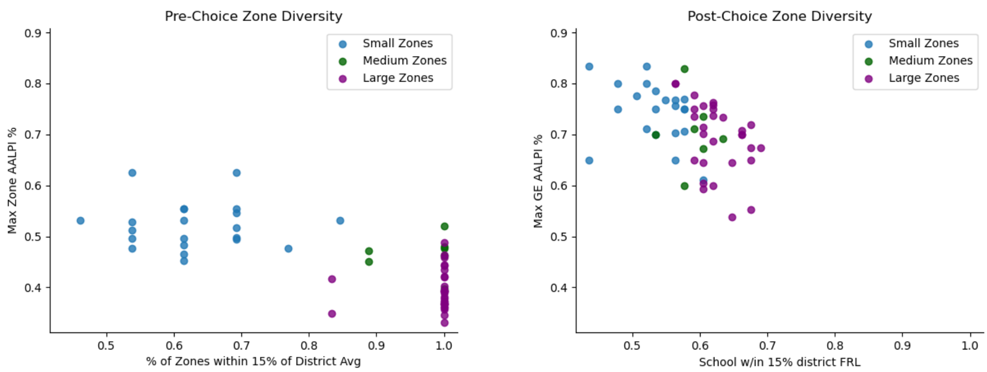

(a) Pre-Choice Diversity Metrics (b) Post-Choice Diversity Metrics

Fig. 5. Figure 5a shows that larger zones perform better on socioeconomic and racial diversity measures. The y axis is the percentage of AALPI students in the zone with the highest percentage, a measure of racial concentration in zones. Figure 5b shows that diverse zones are not necessarily predictive of diverse resulting assignments. The y axis is the percentage of AALPI students in the school with the highest percentage, a measure of racial concentration in schools.

D EQUITY OF ACCESS

To quantify the differing trade-offs between distance and choice for various demographic groups for policies that achieve high diversity, we constructed the following min-weight matching program. Fix a group of students, and create one node for each student. Then, for each program in the district create a number of nodes equal to the capacity of the program, multiplied by the fraction of the total student population that the group of students comprises. Next, add an edge between each student node to each program node, with the weight of the edge equal to the distance between the student and the program. Finally, for input parameters λ and k, add a penalty value λ to the weight of each edge between a student and a program where the student does not rank the program in their top k programs according to the choice model. Then compute the min-weight matching between the student and program nodes, outputting the average distance between students and the programs they were assigned to as well as the fraction of students assigned a program in their top k. By varying the value of λ and computing the min-weight match, a Pareto frontier can be traced of the best achievable distance and choice metrics that a policy can achieve for a given group of students, given that the policy proportionally assigns the students in the group to the programs.

In Figure 6 we show the Pareto frontier that is traced for African American students when $k = 3$, λ varies from 0.8 to 4, and the capacities of each of the programs in the weighted graph were set to be twice as large as would be proportional to the population of African American students. This increase in the program capacities relaxes the constraint that the students be assigned exactly proportionally to the programs. Figure 6 also shows for the average travel distance and fraction receiving one of their top three choices for African American students in the three policies as well as the 2018-2019 assignment. The values of the three policies lie close to the min-weight match line, with Zones + Reserves having the lowest distance and choice, Zones having intermediate values and Priorities having significantly higher distance and choice. The 2018-2019 values lie farther below the min-weight match line than the policies, which is possible because this assignment achieves poorer diversity outcomes.

Fig. 6. Distance and choice trade-off for African American students in the min-weight match, the three policies and the 2018-2019 assignment.

While this program results in a reasonable benchmark, it cannot be directly compared to our other policies for several reasons. First, none of the policies we suggest achieve the level of representative diversity that we require for this program. As a result, we are likely somewhat overestimating the trade-off between choice and distance. Second, this choice policy uses quotas by race, and we cannot use individual student characteristics in student assignment. Consequently, we are likely over-estimating how well we could target outcomes for specific student groups. Regardless, we take this result to suggest that if we require diversity, our policies are operating close to the Pareto frontier with respect to proximity and choice.

E VISUALIZATION TOOL

We developed a metric visualization tool for internal use by district staff and decision-makers. This tool provided geographic, histogram, bar plot, scatter plot, diagonal and zone views for comparing policies.

E.1 Geographic view

Users can visualize a variety of metrics and statistics as they differ across the geography of the city, including by zip code, attendance area, and zone under different zone configurations. These include policy-independent statistics about the various regions of the district (such as the demographic information of students living in each region - see e.g. Figures 1, 7), as well as policy-dependent statistics, including the metrics for optimization under a certain policy calculated for students either living in or assigned to each region. The team used this view to identify geographic trends in the student data in order to better tailor design choices, and to better understand the effect of those choices on student outcomes.

Users can also directly compare the effects of two policies on students living in different parts of the city (Figure 7b). This helps to better contextualize trends of the sort seen in Figure 7a by comparing them to those seen for other policies, including the current assignment system.

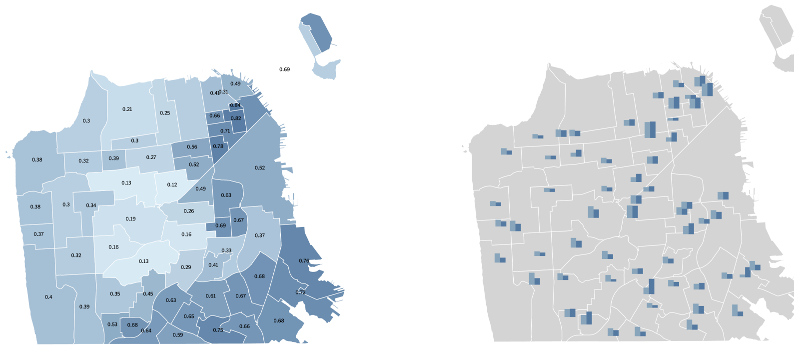

(a) Percentage of students living in each attendance area eligible for FRL

(b) Average distance to assigned school under two different assignment policies

Fig. 7. Geographic view of policy-independent and policy-dependent metrics, averaged over students living in each attendance area.

E.2 Histogram view

For a chosen zone configuration, users can visualize various distribution-based metrics and statistics across each zone. These include both statistics about the schools located in each zone (Figure 8) and the students living in each zone. These visualizations helped reveal inequities between zones under a chosen zone configuration.

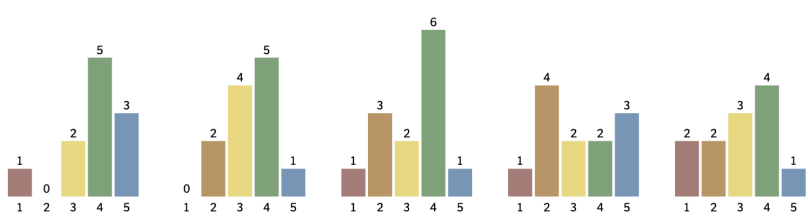

Fig. 8. Count of schools in each zone belonging to each color for English Language test scores (ELA) (as defined by the California School Dashboard) under a certain zone configuration with 5 zones

E.3 Bar plot view

Users can compare policies with the bar plot view, which presents a single metric (such as distance to assigned school) calculated for multiple policies utilizing a single set of parameters (such as whether the CTIP tiebreaker is used). This allows users to quickly identify which policies are most and least effective towards particular goals. Users can also easily adjust the settings to determine the effect of a single parameter—for example, a user can toggle the 'ctip' box to evaluate the effect of using the CTIP1 tiebreaker on outcomes across each policy (see Figure 9).

E.4 Scatter plot view

The scatter plot view illustrates how each policy performs in terms of the trade-offs between two metrics. The user can choose a metric for the x axis and a metric for the y axis, and set the parameters of the simulations; the visualization system will responsively plot each proposal on the graph based on its calculated values for those metrics (see Figure 10). Like the bar plot view, the scatter plot view allows users to easily compare the effects of different parameters on the proposals by quickly toggling each setting.

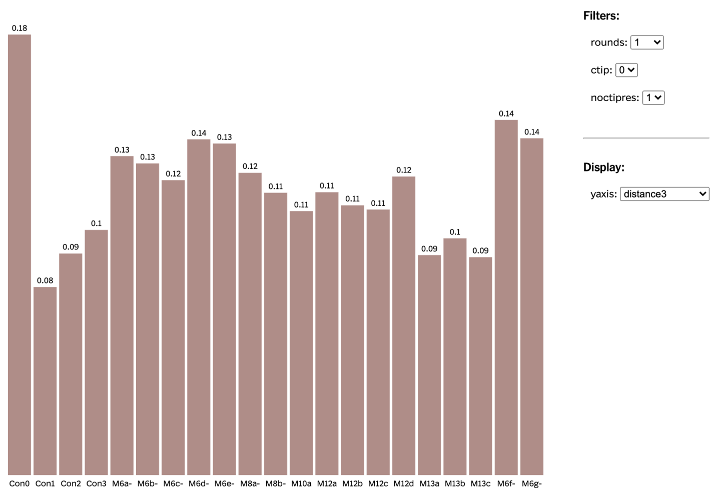

Fig. 9. Proportion of students traveling more than 3 miles to their assigned school under various assignment policies and a single set of parameters

Fig. 10. Plot of average distance vs. percent of students designated under a set of assignment policies and single set of parameters

E.5 Diagonal view

The visualization tool also enables users to further explore the effects of a chosen policy by plotting a metric calculated across students living in each attendance area or zone against the same metric calculated across students assigned to schools in each area or zone (Figure 11). The team used this view to identify whether certain policies effectively disrupted the segregation within the district. If most attendance areas fell near the diagonal (which is marked in the visualization), the demographic patterns under the assignment were similar to those in the district as a whole; policies that created schools highly different from the district's demographic patterns would have more variation across the plot.

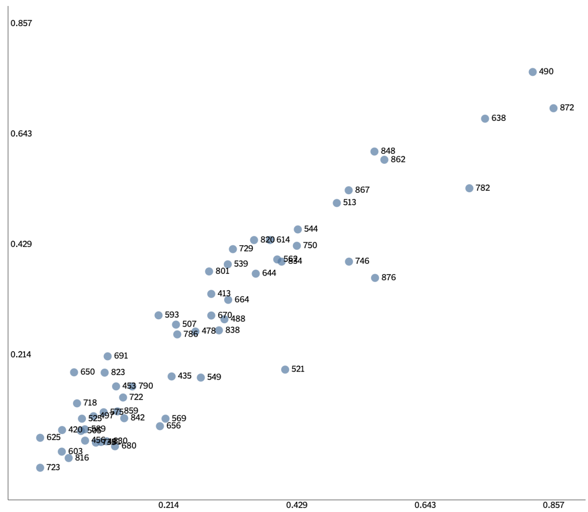

Fig. 11. Percentage of students of Asian ethnicity in each attendance area (x-axis) vs. assigned to schools in each attendance area (y-axis) under a chosen policy

E.6 Zones

Finally, the system helped the team design and understand zone configurations by visualizing zone boundaries across the geography of the city (see,. e.g. Figure 2); such visualizations clarified simple properties of each zone configuration, such as the size and level of contiguity of each zone.