



HAL
open science

Fairness and Diversity in Public Resource Allocation Problems

Nawal Benabbou, Mithun Chakraborty, Yair Zick

► **To cite this version:**

Nawal Benabbou, Mithun Chakraborty, Yair Zick. Fairness and Diversity in Public Resource Allocation Problems. Bulletin of the Technical Committee on Data Engineering, 2019. hal-02308588

HAL Id: hal-02308588

<https://hal.sorbonne-universite.fr/hal-02308588v1>

Submitted on 10 Oct 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fairness and Diversity in Public Resource Allocation Problems*

Nawal Benabbou¹, Mithun Chakraborty², Yair Zick²

¹Sorbonne Université, CNRS, Laboratoire d'Informatique de Paris 6, LIP6 F-75005 Paris, France
nawal.benabbou@lip6.fr

²Department of Computer Science, National University of Singapore, Singapore
{mithun, zick}@comp.nus.edu.sg

Abstract

In this article, we address important extensions to the problem of allocating indivisible items to a population of agents: The agents are partitioned into disjoint groups on the basis of attributes (e.g., ethnicity) and we want the overall utility of the allocation to respect some notion of diversity and/or fairness with respect to these groups. We study two specific incarnations of this general problem. First, we address a constrained optimization problem, inspired by diversity quotas in some real-world allocation problems, where the items are also partitioned into blocks and there is an upper bound on the number of items from each block that can be assigned to agents in each group. We theoretically analyze the price of diversity – a measure of the overall welfare loss due to these capacity constraints – and report experiments based on two real-world data sets (Singapore public housing and Chicago public school admissions) comparing this constrained optimization-based approach with a lottery mechanism with similar quotas. Next, instead of imposing hard constraints, we cast the problem as a variant of fair allocation of indivisible goods – we treat each group of agents as a single entity receiving a bundle of items whose valuation is the maximum total utility of matching agents in that group to items in that bundle; we present algorithms that achieve a standard relaxation of envy-freeness in conjunction with specific efficiency criteria.

1 Introduction

Over the years, the Singapore government has adopted several social integration measures, to accommodate its multi-ethnic and multi-cultural population; one of these is the *Ethnic Integration Policy* (EIP) used by the Housing and Development Board (HDB) since 1989 [27] to determine housing allocations. This government body is charged with the construction of government subsidized public housing estates, and selling them to Singapore residents. The EIP sets upper bounds on the percentage of flats in every estate that can be owned by households of every major ethnic group: since 5 March 2010, every HDB housing block is required to hold no more than 87% Chinese, 25% Malay, and 15% Indian/Others [17, 14]. These ethnic quotas prevent the over-representation of any one group in an estate (resulting in de-facto segregation). HDB uses a lottery to allocate new developments: all applicants who apply for a particular development pick their flats in a random order.

Copyright 0000 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Bulletin of the IEEE Computer Society Technical Committee on Data Engineering

*Parts of this article are published in AAMAS 2018 [7] and accepted to IJCAI 2019 [8].

Under the lottery mechanism, applicants selected later in the order may not get their top choices; in fact, they may be rejected because the quota for their ethnic group has been filled, *even if there are empty flats that they are willing to take*.

Example 1: Consider an estate with 100 flats; Amirah (who is ethnically Malay) receives a queue number of 30 — i.e., she is the 30th person to choose a flat; if, by random chance, 25 ethnic Malays precede her in the queue and pick before her, the ethnic quota for Malays (25%) will be filled and Amirah will no longer be allowed to select a flat in that estate. On the other hand, if Amirah is 110th in line but all families preceding her in the queue are ethnically Chinese, then at most 87 out of the 100 flats will be taken (since the Chinese have an 87% quota) and Amirah will have at least 13 flats to choose from. ♦

Over 72% of Singapore apartments are HDB flats [34], housing an estimated 81% of Singapore residents [18] as of 2018; thus, the HDB public housing market has a significant impact on the life and welfare — both individual and collective — of this island nation [13, 14, 20, 32, 37]. The HDB lottery mechanism, coupled with ethnicity constraints, adds another layer of complexity to what is, at its core, similar to the classic weighted bipartite matching or *assignment* problem [21, 26]: agents (buyer households) have idiosyncratic values/utilities for items (flats) while a central planner (HDB) wishes to allocate at most one item to each agent and vice versa, with the economic criterion of *overall utility/welfare/efficiency* in mind — but also with an added social goal of *promoting diversity*. Inspired by diversity-respecting allocation mechanisms — prevalent not only in Singapore public housing but also in other domains such as matching residents to hospitals in Japan, school admissions in many U.S. cities, and more [19, 36, 16] — we formally study the balance between maximizing allocative welfare and promoting allocative diversity. This underlying problem of allocation/assignment of goods distinguishes our contribution from the extensive literature on fairness and diversity in subset selection (see e.g., [35, 10] and references therein).

In Section 2, we analyze a (simplified) HDB housing market as an extension to the assignment problem where the sets of agents and items are partitioned into subsets called *types* and *blocks* respectively; there is a pre-specified upper bound on the number of agents of each type that can be assigned items in each block, called *type-block capacities*. We analyze the *price of diversity*, i.e., the fractional loss in the overall (optimal) welfare due to these capacity constraints, and relate it to a measure of *type disparity*; we also report simulations based on recent, real-world data sets, comparing our constrained optimization approach with a lottery mechanism with ethnicity quotas in terms of welfare.

In Section 3, we present an alternative approach towards the efficiency-diversity trade-off, drawing on and adding to the rich literature on the fair division of indivisible goods (see e.g., [28, 12, 22, 4, 3]). Here, each type is represented by a *super-agent* that is allocated a *bundle* of items; the super-agent’s valuation of a bundle is given by the maximum utility assignment of items in the bundle to agents of that type. When agent-item utilities are binary, we provide a polynomial-time algorithm that computes an allocation with the maximum (utilitarian) social welfare while satisfying a popular fairness criterion: *envy-freeness up to one item* (EF1) [11]; for arbitrary real-valued utilities, we show experimentally that a heuristic extension of the classic *envy-graph algorithm* due to Lipton et al. [24] produces an EF1 allocation with low *waste* (a new inefficiency concept that we have introduced for this setting).

2 Diversity through hard capacity constraints

First, let us rigorously formulate the welfare maximization problem in a bipartite matching setting under type-block capacity constraints. Detailed proofs of all theoretical results in this section can be found in Benabbou et al. [7]. Throughout the paper, for $s \in \mathbb{N}$, we denote the set $\{1, 2, \dots, s\}$ by $[s]$.

Definition 1 (ASSIGNTC): An instance of the Assignment with Type Constraints (ASSIGNTC) problem is given by: (i) a set N of n agents partitioned into k types N_1, \dots, N_k , (ii) a set M of m items/goods partitioned

into l blocks M_1, \dots, M_l , (iv) a utility $u(i, j) \in \mathbb{R}_+$ for each agent $i \in N$ and each item $j \in M$, and (iv) a capacity $\lambda_{pq} \in \mathbb{N}$ for all $(p, q) \in [k] \times [l]$; λ_{pq} is an upper bound on the number of agents of type N_p allowed in block M_q . W.l.o.g. we assume that for all p, q , $\lambda_{pq} \leq |M_q|$.

An assignment of items to agents, which we will also sometimes call an (N, M) -matching, can be represented by a $(0, 1)$ -matrix $X = (x_{ij})_{n \times m}$ where $x_{ij} = 1$ if and only if item j is assigned to agent i . Our objective is to maximize the utilitarian social welfare or USW, i.e., the sum of agent utilities: $u(X) \triangleq \sum_{i \in N} \sum_{j \in M} x_{ij} u(i, j)$. Clearly, this optimization problem can be formulated as an integer linear program (ILP) as in Figure 1. Here, the first set of inequalities captures our *type-block constraints* while the last three sets are the usual *matching constraints* jointly ensuring that each item (resp. agent) is assigned to at most one agent (resp. item). In general, the decision version of the ASSIGNTC problem is NP-complete (Benabbou et al. [7, Theorem 3.2]) but admits a polynomial-time $\frac{1}{2}$ -approximation algorithm (Benabbou et al. [7, Theorem 3.4]); moreover, the problem can be solved in polynomial(n, m) time by a minimum-cost network flow-based algorithm (Benabbou et al. [7, Theorem 3.6]) if the utility matrix satisfies one of the following two conditions: (i) *type-uniformity* i.e., all agents of a type have identical utilities (for all $p \in [k]$ and for all $j \in M$, there exists $U_{pj} \in \mathbb{R}_+$ such that $u(i, j) = U_{pj}$ for all $i \in N_p$); (ii) *block-uniformity* i.e., all items in a block are clones of each other (for all $q \in [l]$ and for all $i \in N$, there exists $U_{iq} \in \mathbb{R}_+$ such that $u(i, j) = U_{iq}$ for all $j \in M_q$).

We are mainly interested in how the imposition of type-block capacities impacts allocative efficiency. If we denote the set of all valid item assignments X by \mathcal{X} , and all assignments additionally satisfying our type-block constraints by \mathcal{X}_C , then the unconstrained and constrained optimal social welfares for any given utility matrix $(u(i, j))_{n \times m}$ are given by $\text{OPT}(u) \triangleq \max_{X \in \mathcal{X}} u(X)$ and $\text{OPT}_C(u) \triangleq \max_{X \in \mathcal{X}_C} u(X)$. Clearly, $\text{OPT}_C(u) \leq \text{OPT}(u)$ since $\mathcal{X}_C \subseteq \mathcal{X}$. This leads to a natural measure of welfare loss for the ASSIGNTC problem; we call this the *Price of Diversity* (a similar definition appears in [1, 10]): $\text{PoD}(u) \triangleq \text{OPT}(u)/\text{OPT}_C(u)$. By definition, $\text{PoD}(u) \geq 1$, and its exact value depends on agent utilities, but we can bound it, regardless of the utility model, in terms of the *fractional type-block capacities* defined as $\alpha_{pq} \triangleq \lambda_{pq}/|M_q|$ for each $(p, q) \in [k] \times [l]$.

$$\begin{array}{ll}
 \max & \sum_{i \in N} \sum_{j \in M} x_{ij} u(i, j) \\
 \text{s.t.} & \sum_{i \in N_p} \sum_{j \in M_q} x_{ij} \leq \lambda_{pq} \quad \forall p \in [k], \forall q \in [l] \\
 & \sum_{j \in M} x_{ij} \leq 1 \quad \forall i \in N \\
 & \sum_{i \in N} x_{ij} \leq 1 \quad \forall j \in M \\
 & x_{ij} \in \{0, 1\} \quad \forall i \in N, \forall j \in M
 \end{array}$$

Figure 1: ILP formulation of ASSIGNTC.

Theorem 2: For any ASSIGNTC instance, $\text{PoD}(u) \leq 1/\min_{(p,q) \in [k] \times [l]} \alpha_{pq}$.

The following example shows that there is an ASSIGNTC instance whose *PoD* reaches the upper bound in Theorem 2; in other words, this bound is tight.

Example 2: Suppose, $|N_{p_0}| \geq |M_{q_0}|$ for some type-block pair (p_0, q_0) in the set $\text{argmin}_{(p,q) \in [k] \times [l]} \alpha_{pq}$ in an ASSIGNTC instance, and the utilities are given by $u(i, j) = 1$ if $i \in N_{p_0}$ and $j \in M_{q_0}$, $u(i, j) = 0$ otherwise. An optimal unconstrained assignment fully allocates the items in block M_{q_0} to agents in N_{p_0} for a total utility of $|M_{q_0}|$ whereas any optimal constrained assignment allocates exactly $\lambda_{p_0 q_0}$ items in M_{q_0} to agents in N_{p_0} for a total utility of $\lambda_{p_0 q_0}$. Hence, for this family of instances, $\text{PoD}(u) = |M_{q_0}|/\lambda_{p_0 q_0} = 1/\alpha_{p_0 q_0}$. ♦

In general, the bound in Theorem 2 is linear in the number of items i.e., the welfare loss due to hard diversity constraints can be significant in some instances (e.g., Example 2). However, type-block capacities are determined by a central planner in our model; a natural way of setting them is to fix the fractional capacities α_{pq} in advance, and then compute $\lambda_{pq} = \alpha_{pq} \times |M_q|$ when block sizes become available: by committing to a fixed minimum type-block quota α^* (i.e., $\alpha_{pq} \geq \alpha^*$ for all $(p, q) \in [k] \times [l]$), the planner can ensure a $\text{PoD}(u)$ of at most $1/\alpha^*$, regardless of the problem size and utility function. Higher values of α^* reduce the upper bound

on $PoD(u)$ but also increase the capacity of a block for every ethnicity, having an adverse effect on the diversity of block composition – α^* thus functions as a tunable tradeoff parameter between ethnic integration and worst-case welfare loss. In the Singapore housing allocation problem, the EIP fixes a universal percentage cap, slightly higher than the actual corresponding population proportion for every ethnicity. Plugging the current EIP percentages mentioned in Section 1 into the bound in Theorem 2, we get that the Singapore housing system has $PoD(u) \leq \frac{1}{0.87+0.25+0.15} \approx 6.67$. We mention that the price of diversity compares the *best* constrained allocation with the best unconstrained allocation; mechanisms deployed in practice do not necessarily try to find the optimal allocation. For example, the HDB mechanism uses a lottery to allocate flats, and thus may theoretically exhibit greater welfare loss than the bound set in Theorem 2.

Theorem 2 offers a worst-case tight bound on the price of diversity, making no assumptions on agent utilities, but Example 2 suggests that this upper bound is attained when social welfare is solely extracted from a single agent type and a single block. Intuitively, we can improve our bound if a less ‘disparate’ optimal assignment exists. To formalize this notion, we introduce a new parameter. For any optimal unconstrained assignment $X^* \in \mathcal{X}$, let $\beta_p(X^*)$ denote the ratio of the average utility of agents in N_p to the average utility of all agents under X^* . The *inter-type disparity parameter* $\beta(X^*)$ is defined as: $\beta(X^*) \triangleq \min_{p \in [k]} \beta_p(X^*) = \min_{p \in [k]} \frac{u_p(X^*)/|N_p|}{u(X^*)/n}$. Notice that $\beta(X^*)$ is in $(0, 1]$, can be computed in polynomial time and is fully independent of the type-block capacities (it uses X^* , an *unconstrained* optimal assignment). The closer $\beta(X^*)$ is to 1, the lower the disparity between average agents of different types under X^* .

Theorem 3: For any ASSIGNTC instance and any unconstrained optimal assignment $X^* \in \mathcal{X}$, we have $PoD(u) \leq \frac{1/\beta(X^*)}{\sum_{p \in [k]} \nu_p \min_{q \in [l]} \alpha_{pq}}$, where $\nu_p = \frac{|N_p|}{n}$ is the proportion of type $p \in [k]$ in the agent population.

For the Singapore public housing problem, if we use the ethnic proportions reported in the 2010 census report [33] i.e., $|N_1|/n = 0.741$ (Chinese), $|N_2|/n = 0.134$ (Malay), and $|N_3|/n = 0.125$ (Indian/Others) and the same block quotas α_{pq} as before, then in the case of no disparity (i.e., $\beta(X^*) = 1$), a simple calculation based on Theorem 3 shows that $PoD(u) \leq 1.43$ (approx.). Combining Theorems 2 and 3, if we plot the $PoD(u)$ against the disparity parameter $\beta(X^*)$ based on Singapore data, the point corresponding to any ASSIGNTC instance must lie in the shaded region of Figure 2.

2.1 Experimental Analysis




Figure 2: PoD vs disparity parameter for the HDB problem with current EIP quotas and ethnic proportions from Census 2010.

In this section, we present simulations of the ASSIGNTC problem using recent, publicly available datasets: Singapore demographic and housing allocation statistics, and the Chicago public school admission data. We compare the welfare of three assignment mechanisms: unconstrained optimal (maximizing welfare while ignoring the diversity constraints), constrained optimal (finding the optimal allocation under diversity constraints), and one-shot lottery-based (running a lottery with diversity constraints, as is the case for the HDB mechanism).¹ Conducting large-scale surveys to elicit agent preferences over items was beyond the scope of this work, so we simulated utilities based on reasonable models for both problems. We solved both the unconstrained and constrained social welfare maximizations using the Gurobi Optimizer. We refer the reader to <https://git.io/fNhhm> for full implementation details.

¹We generated a uniformly random sequence over agents and assigned to each an item for which it has the highest utility among unassigned items in blocks for which that agent’s ethnicity quota had not been filled yet. We abstract away other complications of actual lottery-based approaches used in our problem domains to focus on diversity constraints.

The Singapore Public Housing Allocation Problem. We collected data on the locations and numbers of flats of recent HDB housing development projects advertised over the second and third quarters of 2017.² Each development constitutes a block in our simulations, for a total of $m = 1350$ flats partitioned into $l = 9$ blocks M_1, \dots, M_9 (Figure 3(a)), consisting of 128, 162, 156, 249, 108, 94, 104, 190, 159 flats respectively. There are pre-specified *categories* of flats, viz. 2-room flexi, 3-room, 4-room, and 5-room; our data set includes lower and upper bounds, $LB(t, q)$ and $UB(t, q)$ respectively, on the monthly cost (loan) for a flat of category t in block M_q for every t and q . We simulate 2 pools of n applicants whose ethnic composition follows the 2010 Singapore census report [33], as shown in Table 1. From the same census report, we collected the average salary $S(p)$ of each ethnicity group $p \in [k]$, given in Singapore dollars: $S(1) = S\$7,326$, $S(2) = S\$4,575$ and $S(3) = S\$7,664$. From publicly available data³ on Singapore’s Master Plan 2014,⁴ we glean the locations of the geographic centers of the 55 planning areas that Singapore is divided into; we also obtained the population sizes of the three ethnicity groups under consideration in each planning area from the General Household Survey 2015 data available from the Department of Statistics, Singapore.⁵ Our block capacities follow latest HDB block quotas [14]: $\alpha_{1q} = 0.87, \alpha_{2q} = 0.25, \alpha_{3q} = 0.15$ for every block M_q .

We simulate 4 utility models; each has one parameter that does not come from the data: **(i) Distance-based** ($Dist(\sigma^2)$): Each agent $i \in N$ has a preferred geographic location $\vec{a}_i \in \mathbb{R}^2$ (chosen uniformly at random within the physical landmass of Singapore) that she would like to live as close as possible to (say, the location of her parents’ apartment, workplace, or preferred school). For every block M_q , we generate the utility of that agent i for apartment $j \in M_q$ by first drawing a sample from the normal distribution $\mathcal{N}(1/d(\vec{a}_i, loc(M_q)), \sigma^2)$, where $loc(M_q) \in \mathbb{R}^2$ is the geographical location of block M_q and $d(\cdot, \cdot)$ represents Euclidean distance, and then renormalizing to make the sum of utilities of each agent for all apartments in M equal to 1. **(ii) Type-based** ($Type(\sigma^2)$): We assume that all agents of the same type (i.e., ethnic group) have the same preferred location (i.e., $\forall p \in [k], \forall i, i' \in N_p, \vec{a}_i = \vec{a}_{i'}$); the rest is similar to the above distance-based model. **(iii) Project approval-based** ($Project(\rho)$): We construct, for each type, a categorical distribution over the 55 planning areas of Singapore, the probability of each area being proportional to the fraction of the sub-population of that type living in that area; for each agent i , we sample a preferred planning area from the above distribution corresponding to i ’s type; if a project M_q is within a radius ρ of the geographic center of agent i ’s preferred planning area, then i approves of the project i.e., $u(i, j) = 1 \forall j \in M_q$, else i disapproves of the project i.e., $u(i, j) = 0 \forall j \in M_q$. **(iv) Price-based** ($Price(\sigma^2)$): Each agent $i \in N_p$ has a salary s_i that is generated according to the normal distribution $\mathcal{N}(S(p), \sigma^2)$. Each flat $j \in M_q$ of category t has a monthly cost p_j chosen uniformly in $[LB(t, q), UB(t, q)]$. The utility that agent i derives from flat j is then defined by $u(i, j) = 1/(p_j - \frac{s_i}{3})^2$, assuming that agent i is willing to pay one-third of her monthly salary on mortgage installments;⁶ the rationale for the utility formula is that a high cost relative to the budget makes flats unaffordable, while a much lower cost indicates unsatisfactory quality, making the agent unhappy in both scenarios.

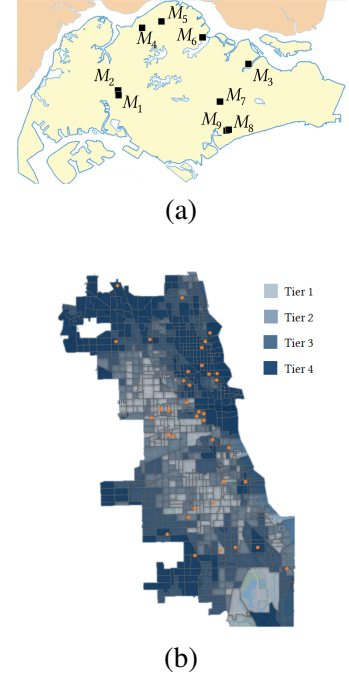


Figure 3: (a) Singapore public housing block locations. (b) Tier statuses of Chicago census tracts and magnet school locations (orange dots).

²<http://www.hdb.gov.sg/cs/infoweb/residential/buying-a-flat/new/bto-sbf>

³<https://data.gov.sg/dataset/master-plan-2014-planning-area-boundary-web>

⁴<https://www.ura.gov.sg/Corporate/Planning/Master-Plan/>

⁵<https://www.singstat.gov.sg/publications/ghs/ghs2015content>

⁶Inspired by the Singapore Central Provident Fund Board-endorsed “3-3-5 rule”, as of 21 Sep 2017.

n	$ N_1 $	$ N_2 $	$ N_3 $
1350	1000	180	170
3000	2223	402	375

Table 1: #applicants (types 1, 2, and 3 are Chinese, Malay, Indian/Others respectively)

For each of our treatments (Figures 4–6), we plot the realized $PoD(u)$ (hatched bar), the theoretical upper bound on $PoD(u)$ as per Theorem 3 (dark gray bar), and the relative loss of the HDB lottery mechanism (i.e., the ratio of $OPT(u)$ to the total utility of the assignment produced by the lottery mechanism) averaged over 100 agent permutations (light gray bar) against the values of the corresponding model parameters (σ^2 or ρ). In order to compare $Dist(\sigma^2)$ with $Type(\sigma^2)$, we vary both σ^2 in $\{1, 5, 10\}$ and n in $\{1350, 3000\}$; the results reported in Figures 4 are on average performance over 100 randomly generated instances. Our first observation is that, in all our experiments, $Dist(\sigma^2)$ exhibits virtually no utility loss due to the imposition of type-block constraints (see the hatched bars in Figures 4(a)). This is because utilities in $Dist(\sigma^2)$ are independent of ethnicity, resulting in a very low value for the inter-type disparity parameter β (indicated by the dark gray bars) — in fact, for any utility model where utilities are independent of ethnicity, the expected value of the disparity parameter is 1. For the $Type(\sigma^2)$ model-based utilities, the disparity parameter is somewhat higher (utilities do strongly depend on ethnicity), resulting in a higher $PoD(u)$ (see the hatched bars in Figures 4(b)). Despite making no attempt to optimize social welfare

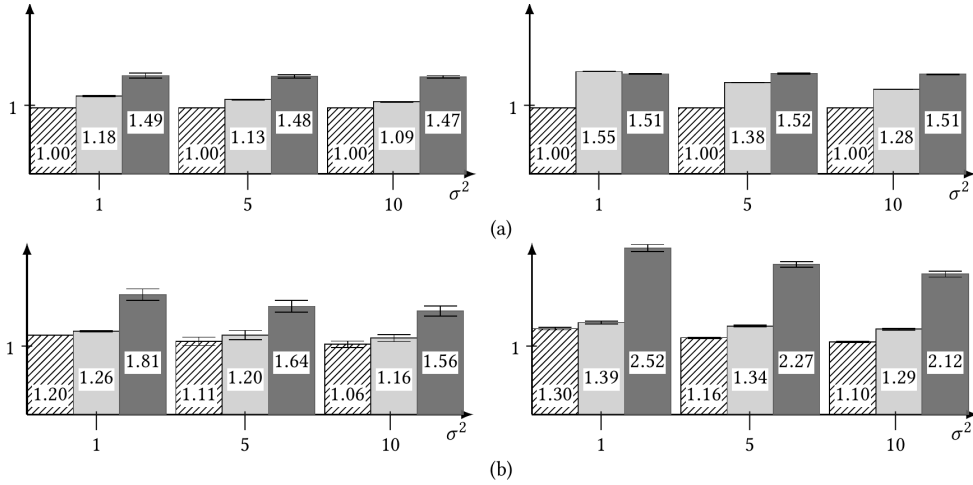


Figure 4: Averaged utility losses for (a) $Dist(\sigma^2)$ and (b) $Type(\sigma^2)$ with $n = m = 1350$ (left) and $n = 3000, m = 1350$ (right).

under type-block constraints, the HDB lottery mechanism does surprisingly well when the number of agents equals the number of apartments (see the light gray bars in the left part of Figure 4), extracting at least 84% of the optimal unconstrained welfare under the $Dist(\sigma^2)$ utility model, and at least 79% of the social welfare under the $Type(\sigma^2)$ model. However, the lottery-induced welfare is negatively impacted by the number of agents (see the light gray bars in the right part of Figure 4); for instance, it only extracts 65% of the optimal unconstrained welfare under $Dist(1)$ with $n = 3000$ and, in fact, the lottery-induced welfare loss for this treatment even exceeds the theoretical upper bound on the price of diversity.

For $Project(\rho)$, we use the fact that one degree of latitude or longitude at the location of Singapore corresponds to roughly 111 km to compute distances; we vary ρ in $\{5, 7.5, 10\}$ (in km). The results averaged over 100 runs are provided in Figure 5. In all instances, $PoD(u)$ is almost one and the lottery-induced welfare is also nearly as good, achieving at least 87% of the unconstrained optimum for 1350 agents and practically 100% for 3000 agents; the disparity parameter is also consistently close to its ideal value of 1, keeping the upper bound at around 1.45 regardless of the radius. Thus, this can be considered an example of a utility model for which the lottery mechanism virtually implements a constrained optimal allocation for a wide range of model

parameters. Finally, for $Price(\sigma^2)$, we vary σ^2 in $\{0, 10, 50\}$; the results obtained by averaging over 100 runs




Figure 5: Averaged utility losses for $Project(\rho)$ with $n = m = 1350$ (left) and $n = 3000, m = 1350$ (right).

are given in Figure 6. While the price of diversity is practically equal to one in all instances, the welfare loss observed with the lottery mechanism drastically increases with σ^2 (recall that agents from the same ethnicity group have identical preferences when $\sigma^2 = 0$): for instance, for 1350 agents, it extracts 98% of the optimal unconstrained welfare under $Price(0)$ while it only extracts 35% of this value under $Price(50)$. These numerical tests show that utility models exist for which the lottery mechanism may perform poorly compared to the optimal constrained allocation mechanism, even in allocation problems with a very low price of diversity.




Figure 6: Averaged utility losses for $Price(\sigma^2)$ with $n = m = 1350$ (left) and $n = 3000, m = 1350$ (right).

Chicago Public School Admissions. Chicago Public Schools (CPS) is one of the largest school districts in the U.S.A.,⁷ overseeing more than 600 schools of various types.⁸ The application and selection processes for these schools involve a number of computerized lotteries, with a significant number of entry-level seats in magnet and selective enrollment schools being filled by lotteries based on a *tier system* based on the family *socio-economic status* (SES) as part of a social integration policy. The city computes a multi-factor, composite SES score for each of the census tracts that Chicago is divided into, and places each tract in one of four tiers based on its score in such a way that each tier contains (roughly) a quarter of school-aged children. The tier of a child is determined by their residential address. Of the seats in each school earmarked for a *citywide SES lottery* or *general lottery*, an equal number is allocated to each tier, and there is an upper limit on the number of schools that a child may apply to (see [30, 29]). We apply our setup to a simplified version of the CPS student-seat allocation scenario.

We collected data from the Chicago Public Schools website⁹ on the locations of magnet schools in Chicago (which use a lottery mechanism to select students), as well as the total number of students enrolled in these

⁷http://www.cps.edu/About_CPS/At-a-glance/Pages/Stats_and_facts.aspx

⁸<http://cpstiers.opencityapps.org/about.html>

⁹<http://cps.edu/Pages/home.aspx>

schools in 2018, which we divided by 9 to obtain the approximate number of first-graders (there are nine grades in total). This leads us to instances with $l = 37$ blocks (schools) and $m = 2261$ items (available seats) in total. In this school admission problem, students are partitioned into $k = 4$ types, viz. Tiers 1, 2, 3 and 4, depending on their residence (see Figure 3(b)¹⁰). In our experiments, we simulate 2 pools of n students where the type composition follows the real-world proportion, as shown in Table 2. Our type-block capacities are $\lambda_{pq} = 0.25|M_q|$ for every pair (p, q) . For our student-seat utility simulations, we use the distance-based utility model $Dist(\sigma^2)$ we introduced in the housing domain, with the following important modifications: we choose the preferred location of a student uniformly at random from the collection of census tracts (polygons) belonging to her tier (see Figure 3(b)), where the position of every polygon is approximated by taking the averaged coordinates of its extreme points; we reset each student’s utility to 0 for any school ranked 21st or lower in the preference ordering induced by her utilities (since students are allowed to apply to at most 20 schools), and then renormalize the utilities.

n	$ N_1 $	$ N_2 $	$ N_3 $	$ N_4 $
2261	613	622	533	493
5000	1355	1375	1180	1090

Table 2: #students (type p is Tier p for each $p \in [4]$.)

In our experiments, we vary σ^2 in $\{0, 10, 50\}$, and report 100-run averages of the same measurements as in the Singapore-based simulations (Figure 7). We observe that both the price of diversity (hatched) and the loss of the lottery mechanism (light gray bar) decrease as σ^2 increases, remaining well below the Theorem 3 bound (dark gray). However, the lottery mechanism loss is quite high in all instances and, just as in the Singapore case, gets worse for a higher number of

students. Our experiments suggest that the lottery mechanism is better suited to problems with an equal number of agents and items.




Figure 7: Averaged utility losses for $Dist(\sigma^2)$ with $n = m = 2261$ (left) and $n = 5000, m = 2261$ (right) in our Chicago-based simulations.

3 Group Fairness in Allocation

Up to this point, we have only explored the welfare loss due to capacity constraints; however, allocative efficiency is only one facet of group fairness. In some settings, groups might receive an overall worse outcome from the allocation mechanism, as compared to others. This is known in the fairness literature as *envy*. In this

¹⁰Based on data from <http://cpstiers.opencityapps.org/> and <http://cps.edu/ScriptLibrary/Map-SchoolLocator/index.html>.

section, we explore how envy-freeness (i.e., having no agent group envy another’s allocation) affects the allocation outcome. We work with a model similar to that in Section 2 with two differences: (i) the items M are not partitioned into blocks (or, equivalently, there is one block); (ii) we assume that for each agent $i \in N$ (resp. item $j \in M$), there is at least one item $j \in M$ (resp. agent $i \in N$) with $u(i, j) > 0$. We adopt an alternative view of diversity-respecting assignment as a task of allocating bundles (i.e., disjoint subsets of M) to *super-agents* (i.e., the types N_1, \dots, N_k) in a manner that is both *fair* and *efficient* (see [8] for further details and complete proofs).

Each type p is represented by a *super-agent*. We define $v_p(S)$, the *valuation* of any super-agent $p \in [k]$ for any bundle of items $S \subseteq M$, as the maximum utilitarian social welfare of matching items in S to agents of type N_p ; $v_p(\cdot)$ is a monotonic, submodular set function (see [8, Theorem 1]). Moreover, our model does not satisfy the additive bundle-valuation assumption or the public goods assumption prevalent in most prior work (see e.g., [12, 5, 2, 31] and references therein), necessitating novel solution techniques.

Definition 4 (Allocation): An allocation \mathcal{A} is a collection of bundles $M_1^{\mathcal{A}}, \dots, M_k^{\mathcal{A}}$, such that $M_1^{\mathcal{A}} \cup \dots \cup M_k^{\mathcal{A}} \subseteq M$ and $M_p^{\mathcal{A}} \cap M_q^{\mathcal{A}} = \emptyset$ for all $p, q \in [k]$ with $p \neq q$, along with a maximum-USW matching between each type N_p and its allocated bundle $M_p^{\mathcal{A}}$ for all $p \in [k]$, thereby inducing a unique (N, M) -matching $X^{\mathcal{A}} = (x_{ij}^{\mathcal{A}})_{i \in N, j \in M}$.

We call $M_0^{\mathcal{A}} = M \setminus \bigcup_{p \in [k]} M_p^{\mathcal{A}}$ the set of *withheld items* under allocation \mathcal{A} . In general, withholding items means that an allocation violates, by definition, the *completeness* condition, a commonly used efficiency criterion. A type p ’s *marginal utility* for an item j is the difference in p ’s valuation of S with and without item j : $\Delta_p(S; j) \triangleq v_p(S \cup \{j\}) - v_p(S)$ if $j \notin S$; $\Delta_p(S; j) \triangleq v_p(S) - v_p(S \setminus \{j\})$ if $j \in S$. We say that an item $j \in M_p^{\mathcal{A}}$ is *unused* under \mathcal{A} if it is either unassigned in the corresponding matching between N_p and $M_p^{\mathcal{A}}$ or is assigned to an agent $i \in N_p$ such that $u(i, j) = 0$. *Cleaning* is the procedure of revoking all unused items from an allocation and putting them in the withheld set. An item $j \in M$ is *wasted* by an allocation \mathcal{A} if it is either withheld (i.e., $j \in M_0^{\mathcal{A}}$) or allocated to a type p and unused, although there is some other type q with $\Delta_q(M_q^{\mathcal{A}}; j) > 0$. A *non-wasteful* allocation has no wasted items. Non-wastefulness is a reasonable efficiency concept in this setting; in fact, it turns out to be a relaxation ([8, Lemma 1]) of a popular efficiency criterion, *Pareto optimality*: an allocation is Pareto optimal among types if the realized bundle-value of no type under this allocation can be strictly improved without strictly diminishing that of another.

We base our fairness criterion on the concept of *envy*: a type p envies a type q if $v_p(M_p^{\mathcal{A}}) < v_p(M_q^{\mathcal{A}})$; p envies q up to ν items, $\nu \in [|M_q^{\mathcal{A}}|]$, if there is a subset $C \subseteq M_q^{\mathcal{A}}$ of size $|C| = \nu$ such that $v_p(M_p^{\mathcal{A}}) \geq v_p(M_q^{\mathcal{A}} \setminus C)$ and, for every subset $C' \subseteq M_q^{\mathcal{A}}$ with $|C'| < \nu$, $v_p(M_p^{\mathcal{A}}) < v_p(M_q^{\mathcal{A}} \setminus C')$. Ideally, we want no type to envy another but such an allocation may not exist; a relaxation that always exists is the following:

Definition 5 (Envy-freeness up to one item [11]): Allocation \mathcal{A} is *envy-free up to one item* (EF1) among types if for any two types $p, q \in [k]$, p either does not envy q or envies q up to one item i.e., there exists an item $j \in M_q^{\mathcal{A}}$ such that $v_p(M_p^{\mathcal{A}}) \geq v_p(M_q^{\mathcal{A}} \setminus \{j\})$.

We want our allocation to be not just EF1 among types (thereby respecting diversity) but also efficient in one of the ways discussed above. Our first result in this vein applies to the *binary utility model*: $u(i, j) \in \{0, 1\}$, $\forall i \in N, \forall j \in M$. This captures scenarios where each agent either approves or disapproves of an item but does not distinguish among its approved items. Moreover, in formal conversations with stakeholders, we have found that a binary utility model is likely consistent with how agents value items in many real-world situations, e.g., in housing markets, a potential buyer might want a flat of a particular category only (such as a 3BHK within a 5-km radius of her workplace), being indifferent among flats of the same category.

Theorem 6: For any problem instance with a binary utility model, Algorithm 1 computes in $poly(n, m)$ time an EF1 allocation that also maximizes the utilitarian social welfare of the induced (N, M) -matching.

It is easy to see that optimal utilitarian social welfare automatically implies Pareto optimality among types, and hence non-wastefulness. Thus, Algorithm 1 solves the fair and efficient allocation problem for binary utilities.

Algorithm 1: Maximum-size Matching with Envy-Induced Transfers

- 1 Compute a maximum-size matching of bipartite graph (N, M) such that there is an edge between i and j iff $u(i, j) = 1$, and clean the resulting allocation; designate the subset of items matched to agents in N_p as type p 's allocated bundle $M_p^A \forall p \in [k]$.
 - 2 **/*Envy-Induced Transfers*/**
 - 3 **while** there are two types p, q such that p envies q up to more than 1 item. **do**
 - 4 Find item $j' \in M_q^A$ such that $\Delta_p(M_p^A; j') > 0$.
 - 5 $M_q^A \leftarrow M_q^A \setminus \{j'\}; M_p^A \leftarrow M_p^A \cup \{j'\}$.
 - 6 Compute a maximum-size (N_p, M_p) -matching.
 - 7 **end**
-

For more general utilities in \mathbb{R}_+ , an algorithm that guarantees a similarly fair and efficient allocation remains elusive. However, we note that it is possible to obtain a type-complete TEF1 allocation in polynomial time by a natural extension (called **L** hereafter) of the algorithm due to Lipton et al. [24]: iterate over the items $j \in M$, giving item j to a type, say p , that is currently not envied by any other type for its current bundle M_p ; compute an optimal matching with the augmented bundle $M_p \cup \{j\}$; construct the *envy graph* where there is a directed edge from a type q to a type r whenever q envies r and eliminate any cycle in this graph by transferring the bundle of every type on this cycle to its predecessor on this cycle (to ensure that there is an unenvied type in each iteration), followed by re-matching within each such type. Although no item is withheld, it is possible for the final allocation to be wasteful: an item may be allocated to a type which has zero marginal utility for it or may *become* unassigned after a bundle is transferred between types.

One heuristic that could reduce waste is the following: in each iteration, find an item-type pair having the maximum marginal utility among all currently unenvied types and all unallocated items (breaking further ties uniformly at random, say), and allocate that item to that type. We call **L**, augmented with this heuristic, **H**.

To test how this marginal utility maximization heuristic performs in practice, we experimentally compared procedures **L** and **H** using the percentage of items wasted averaged over runs, denoted by *%Waste*, as our performance metric. We simulated two data sets with $n = 100$ agents partitioned into $k = 3$ types: UNEQUAL (ethnic proportions following Singapore 2010 census [33]): $|N_1| = 74, |N_2| = 13, |N_3| = 13$; EQUAL: $|N_p| \approx n/k$ for all types $p \in [k]$. For each agent, we sampled m numbers uniformly at random from $[0, 1]$ and

Data set	#Items	%Waste	
		L	H
UNEQUAL	50	13%	0
	100	39%	0
EQUAL	50	0%	0
	100	0.005%	0

normalized them to generate utilities for all m items, with $m \in \{50, 100\}$. The results are shown in the adjoining table: the main observation is that **H** produces *zero* waste for *all* experiments. Thus, augmenting **L** with a simple heuristic produces a surprising improvement in performance over a wide range of problem parameters.

4 Discussion and future work

One of the extensions of the work presented here that we are currently pursuing is a rigorous analysis of the lottery mechanism with diversity quotas which we experimentally compared with our constrained optimization benchmark in Section 2.1. We are trying to assess whether certain lotteries are better than others in maintaining diverse but efficient outcomes in theory and in practice i.e., how the different parameters (the number of types, their respective percentage caps, sizes, and their utility structures) interact with the randomness of the draws to affect the welfare of the entire population as well as welfare-discrepancies among types.

One other major direction we are investigating is an extension of/alternative to Algorithm 1 for arbitrary real-

valued utilities. Several other possible approaches towards a tradeoff between fairness/diversity and efficiency are also worth exploring: diversity through the optimization of carefully constructed objective functions [23, 1]; extensions of non-envy-based fairness concepts (group-wise egalitarian welfare, maximin shares [3, 6], etc.) to our matching-based setting, and so on.

Acknowledgments

Chakraborty and Zick are supported by Singapore NRF Fellowship R-252-000-750-733, and Benabbou by ANR project 14-CE24-0007-01-Cocorico-CoDec; a major part of the work was done when Benabbou was a post-doctoral research fellow at National University of Singapore (2017-18), supported by Singapore NRF Fellowship R-252-000-750-733. The authors would like to thank Xuan-Vinh Ho, Jakub Sliwinski (supported by MOE grant R-252-000-6255-133), and Edith Elkind as co-authors of publications on which this article is based, and Ayumi Igarashi for insightful discussions. Thanks are also due to the anonymous reviewers of AAMAS 2018 and IJCAI 2019, and the attendees of COMSOC 2018 and FAMAS 2019 where parts of this work were presented.

References

- [1] F. Ahmed, J. P. Dickerson, and M. Fuge. Diverse Weighted Bipartite b -Matching. *Proc. 26th IJCAI*, pp. 35–41, 2017.
- [2] M. Aleksandrov and T. Walsh, Toby. Group envy freeness and group Pareto efficiency in fair division with indivisible items. *Künstliche Intelligenz*, pp. 57–72, 2018.
- [3] S. Barman and S. K. Krishnamurthy. Approximation Algorithms for Maximin Fair Division. *Proc. 18th EC*, pp. 647–664, 2017.
- [4] S. Barman, S. K. Krishnamurthy, and R. Vaish. Finding Fair and Efficient Allocations. *Proc. 19th EC*, pp. 557-574, 2018.
- [5] S. Barman, S. K. Krishnamurthy, R. Vaish. Greedy algorithms for maximizing Nash social welfare. *Proc. 17th AAMAS*, pp. 7–13, 2018.
- [6] S. Barman, A. Biswas, S. K. Krishnamurthy, and Y. Narahari. Groupwise maximin fair allocation of indivisible goods. *Proc. 32nd AAI*, pp. 917–924, 2018.
- [7] N. Benabbou, M. Chakraborty, X. V. Ho, J. Sliwinski, and Y. Zick. Diversity Constraints in Public Housing Allocation. *Proc. 17th AAMAS*, pp. 973–981, 2018.
- [8] N. Benabbou, M. Chakraborty, E. Elkind, and Y. Zick. Fairness Towards Groups of Agents in the Allocation of Indivisible Items *Accepted to 28th IJCAI*, 2019.
- [9] S. Bouveret, Y. Chevaleyre, and N. Maudet. Fair Allocation of Indivisible Goods. *Handbook of Computational Social Choice*, ed. F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, ch. 12, pp. 284–310, Cambridge University Press, 2016.
- [10] R. Bredereck, P. Faliszewski, A. Igarashi, M. Lackner, and P. Skowron. Multiwinner Elections with Diversity Constraints. *Proc. 32nd AAI*, pp. 933–940, 2018.
- [11] E. Budish The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, University of Chicago Press, Vol. 119, No. 6, pp. 1061–1103, 2011.
- [12] I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, J. Wang. The unreasonable fairness of maximum Nash welfare. *Proc. 17th EC*, pp. 305–322, 2016.
- [13] B. Chua. Race relations and public housing policy in Singapore. *Journal of Architectural and Planning Research*, pp. 343–354, 1991.
- [14] Y. Deng, T. F. Sing, and C. Ren The story of Singapore’s public housing: From a nation of home-seekers to a nation of homeowners. *The Future of Public Housing: Ongoing Trends in the East and the West*, ed. J. Chen, M. Stephens, and Y. Man, ch. 7, pp. 103–121, Springer, 2013. ISBN 978-3-642-41621-7.

- [15] B. Fain, A. Goel, and K. Munagala. The core of the participatory budgeting problem. *Proc. 12th WINE*, pp. 384–399, 2016.
- [16] D. Fragiadakis and P. Troyan. Improving matching under hard distributional constraints. *Theoretical Economics*, Wiley Online Library, Vol. 12, No. 2, pp. 863–908, 2017.
- [17] Housing and Development Board, Singapore. Policy changes to support an inclusive and cohesive home [Press release]. 5 Mar 2010. <http://www.nas.gov.sg/archivesonline/speeches/record-details/809e76bf-115d-11e3-83d5-0050568939ad>.
- [18] Housing and Development Board, Singapore. Annual Report 2016/2017: Key Statistics. 2017 <http://www10.hdb.gov.sg/ebook/AR2018/key-statistics.html>.
- [19] Y. Kamada and F. Kojima. Efficient matching under distributional constraints: Theory and applications. *American Economic Review*, Vol. 105, No. 1, pp. 67–99, 2015.
- [20] S. Y. Phang and K. Kim. Singapore’s Housing Policies: 1960–2013. *Frontiers in Development Policy: Innovative Development Case Studies*, pp. 123–153, 2013.
- [21] H. W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics*, 2(1-2):83–97, 1955.
- [22] D. Kurokawa, A. D. Procaccia, and J. Wang. When can the maximin share guarantee be guaranteed? *Proc. 30th AAAI*, pp. 523–529, 2016.
- [23] J. Lang and P. K. Skowron. Multi-Attribute Proportional Representation. *Proc. 30th AAAI*, pp. 530–536, 2016.
- [24] R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. *Proc. 5th EC*, pp. 125–131, 2004.
- [25] L. Lovász and M. D. Plummer. Matching theory. *American Mathematical Society*, Vol. 367, 2009.
- [26] J. Munkres. Algorithms for the assignment and transportation problems. *Journal of SIAM*, Vol. 5, No. 1, pp. 32–38, 1957.
- [27] Parliament of Singapore. Better racial mix in HDB housing estates. Parliament Debates: Official Report. 16 Feb 1989. Vol. 52, cols. 650–668.
- [28] A. D. Procaccia and J. Wang. Fair enough: Guaranteeing approximate maximin shares. *Proc. 15th EC*, pp. 675–692, 2014.
- [29] K. Quick. Chicago Public Schools: Ensuring Diversity in Selective Enrollment and Magnet Schools. *The Century Foundation*. <https://tcf.org/content/report/chicago-public-schools>, 14 Oct 2016.
- [30] Chicago Public Schools. Chicago Public Schools Policy Manual: Admissions Policy for Magnet, Selective Enrollment and Other Options For Knowledge Schools and Programs. Section 202.2. Board Report 17-0426-PO2. Available at <http://policy.cps.edu/Policies.aspx>, 26 Apr 2017.
- [31] E. Segal-Halevi and W. Suksompong. Democratic Fair Allocation of Indivisible Goods. *Proc. 27th IJCAI*, pp. 482–488, 2018.
- [32] L. L. Sim, S. M. Yu, and S. S. Han. Public housing and ethnic integration in Singapore. *Habitat International*, Vol. 27, No. 2, pp. 293–307, 2003.
- [33] Department of Statistics, Singapore. Singapore 2010 Census: Key Indicators of the Resident Population. 2010.
- [34] Department of Statistics, Singapore. Singapore in Figures. 2018.
- [35] J. Stoyanovich, K. Yang, and H.V. Jagadish. Online set selection with fairness and diversity constraints. *Proc. 21st EDBT*, pp. 241–252, 2018.
- [36] U.S. Department of Education, Office of Elementary and Secondary Education. Improving Outcomes for All Students: Strategies and Considerations to Increase Student Diversity. Washington, D.C. <https://www.ed.gov/diversity-opportunity>, 19 Jan 2017.
- [37] M. Wong. Estimating the distortionary effects of ethnic quotas in Singapore using housing transactions. *Journal of Public Economics*, Vol. 115, pp. 131–145, 2014.