

Diversity Constraints in Public Housing Allocation

Nawal Benabbou
National University of Singapore
Singapore
benabbou@comp.nus.edu.sg

Mithun Chakraborty
National University of Singapore
Singapore
dcsmc@nus.edu.sg

Xuan-Vinh Ho
National University of Singapore
Singapore
hxvinh@comp.nus.edu.sg

Jakub Sliwinski*
ETH Zürich
Switzerland
jsliwinski@ethz.ch

Yair Zick
National University of Singapore
Singapore
dcsyaz@nus.edu.sg

ABSTRACT

The state of Singapore operates a national public housing program, accounting for over 70% of its residential real estate. Singapore uses its housing allocation program to promote ethnic diversity in its neighborhoods; it does so by imposing ethnic quotas: every ethnic group must not own more than a certain percentage in a housing project, thus ensuring that every neighborhood contains members from each ethnic group. However, imposing diversity constraints naturally results in some welfare loss. Our work studies the tradeoff between diversity and (utilitarian) social welfare from the perspective of computational economics. We model the problem as an extension of the classic assignment problem, with additional diversity constraints. While the classic assignment program is poly-time computable, we show that adding diversity constraints makes the problem computationally intractable; however, we identify a $\frac{1}{2}$ -approximation algorithm, as well as reasonable agent utility models which admit poly-time algorithms. In addition, we study the *price of diversity*: this is the loss in welfare incurred by imposing diversity constraints; we provide upper bounds on the price of diversity as functions of natural problem parameters. Finally, we use recent, public demographic and real-estate data from Singapore to create a simulated framework testing the welfare loss due to diversity constraints in realistic large-scale scenarios.

KEYWORDS

Optimal Assignment; Type Constraints; Diversity

ACM Reference Format:

Nawal Benabbou, Mithun Chakraborty, Xuan-Vinh Ho, Jakub Sliwinski, and Yair Zick. 2018. Diversity Constraints in Public Housing Allocation. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018*, IFAAMAS, 9 pages.

1 INTRODUCTION

More than 80% of Singapore citizens and permanent residents live in public housing projects [26]; these apartments are sold in a large-scale public market centrally managed by a government body called

*This work was done while this author was at the National University of Singapore.

the Housing and Development Board (HDB). In addition to providing a public good — affordable apartments in a small country with little real estate — HDB estates serve a role in the social integration of Singapore’s diverse ethnic groups (Chinese, Malay and Indian/Others). As per the Ethnic Integration Policy introduced in 1989 [36], every public housing development must hold a certain percentage of every major ethnic group, which is somewhat proportional to the true percentages of these groups in the general Singaporean population; for example, since 5 March 2010, every HDB housing block is required to hold no more than 87% Chinese, 25% Malay, and 15% Indian/Others [18, 24]. Ethnic quotas ensure a diverse population composition at the block-level, preventing the de-facto formation of segregated ethnic communities in public housing estates. HDB uses a lottery mechanism to allocate new developments: all applicants who apply for a particular development pick their flats in a random order; however, these ethnicity constraints introduce some peculiarities. For example, consider an applicant i of Chinese ethnicity applying for an apartment block with 100 flats, up to 87 of which may be assigned to ethnically Chinese applicants, and at most 25 of which can be assigned to ethnically Malay applicants. Assume that i is 90th in line to select an apartment. If at least 87 Chinese applicants were allowed to choose a flat before i , the Chinese ethnic quota for the block will have been filled and applicant i will no longer be eligible for the block. On the other hand, suppose that i is 105th in line to select an apartment; if 40 Malay applicants end up before i in the lottery, then 15 of them will be rejected, and i will have a spot¹.

As the example above shows, ethnic quotas add another layer of complexity to what is, at its foundation, a straightforward allocation problem. Indeed, the allocation of public housing is an economic problem similar to the classic assignment problem: a central planner (HDB) wishes to allocate goods (apartments) to agents (residents) in a manner satisfying certain economic criteria. Diversity, on the other hand, is a social goal external to the underlying economic domain; imposing it may result in reduced social welfare.

1.1 Our Contributions

We study the interplay between diversity and social welfare in the public housing market; we model this as an assignment problem with additional *type-block constraints*: agents are of multiple *types*

¹While this example is, of course, stylized, the effects it describes are quite real: one often hears stories of young couples who arrive at the HDB office to select a flat, only to be notified that their ethnic quota has just been filled.

and goods are divided into *blocks*. A limited number of goods in each block can be allocated to agents of each type; we call these upper bounds *type-block capacities*. These restrictions result in several interesting outcomes. While the optimal assignment problem is well-known to be poly-time solvable [29], we show that imposing type-block constraints makes it computationally intractable (Section 3). However, we identify a $\frac{1}{2}$ -approximation algorithm (Section 3.1), as well as some agent utility models for which one can find optimal assignments with type-block constraints in polynomial time (Section 3.2). Next, we show that the potential utility loss from imposing type-block constraints – which we term the *price of diversity* as in [4] – can be bounded by natural problem parameters (Section 4). Finally, we analyze the empirical price of diversity on simulated instances generated from real HDB data (Section 5).

1.2 Related work

One can think of public housing allocation as a bipartite matching problem [33] where an edge from an agent i to an apartment j is weighted with the utility the agent will receive if she is allocated the apartment. There is a rich literature on weighted bipartite matching problems (also known as assignment problems [34]), and polynomial-time algorithms for the unconstrained version have long been known (e.g. [29]). Several generalizations and/or constrained versions have been studied, e.g. recent work by Lian et al. [2018] who allow each agent (resp. item) to be matched to multiple items (resp. agents) but within upper and lower capacities. The assignment problem with subset constraints studied by Bauer [2004] can be thought of as a special case of our problem, with a single block or a single type. If all agents of each type have identical utilities for all apartments in each block, and each type-block capacity is smaller than both the corresponding type and block sizes, then our problem reduces to a special case of the polynomial-time solvable *capacitated b-matching* on a bipartite graph [5].

Singapore’s public housing is our primary motivating domain but type-block constraints can naturally arise in many other settings [27, 44]. For example, consider the course allocation problem analyzed by Budish and Cantillon [2012]; one might require that each course has students from different departments and impose maximal quotas to ensure this. In public school allocation [1, 3, 37], one might require that certain schools admit students from diverse neighborhoods to prevent de-facto segregation. Other examples include matching medical interns or residents to hospitals [40], allocating subsidized on-campus housing to students [2], appointing teachers at public schools in different regions as done by some non-profit organizations [21], or assigning first year business school students to overseas programs [21]. This line of work mainly explores the interaction between individual selfish-rational behavior and allocative efficiency (e.g. Pareto-optimality) of matching mechanisms, under either ordinal preferences or cardinal utilities, one-sided or two-sided (see, e.g. [6, 7, 11, 12] and references therein); we, on the other hand, focus on the impact of type-block constraints on welfare loss, when agents’ utilities are known to a central planner.

A relevant body of work is that on the fair allocation of indivisible goods (see, e.g., [8, 9, 15, 30, 39]): here, fairness is quantified in terms of agents’ realized utilities or preferences whereas we deal with the proportionate representation of groups in the allocation with

respect to an attribute different from utility. Some recent literature [13, 31] addresses diversity in a subset selection setting; Ahmed et al. [2017] “treat[] diversity as an objective, not a constraint” in a b -matching context and use a ‘soft’ approach towards encouraging diversity, whereas we enforce diversity through hard constraints.

Immorlica et al. [2017] study the efficiency of lottery mechanisms such as the ones used by HDB to allocate apartments; however, their work does not account for ethnicity-block constraints which, as we show here, have a significant effect on allocative efficiency.

1.3 The Singapore Public Housing Allocation System

The Singapore public housing system, managed by the HDB, provides low-cost apartments to Singapore citizens and permanent residents. Public housing is a dominant force in Singapore: as of 2017, approximately 73.3% of apartments in Singapore are HDB flats [20]. New HDB flats are purchased directly from the government, which offers them at a heavily subsidized rate. New apartments are typically released at quarterly sales launches: these normally consist of plans for several estates at various locations around Singapore, an estate comprising four or five blocks (each apartment block has approximately 100 apartments) sharing some communal facilities (e.g. a playground, a food court, a few shops etc.). Estates take between 3 to 5 years to complete, during which HDB publicly advertises calls to ballot for an apartment in the new estate. A household (say, a newly married couple looking for a new house) usually ballots for a few estates (balloting is cheap: only S\$10 per application [25]). HDB allocates apartments using a lottery: all applicants to a certain estate choose their flat in some random order; they are only allowed to select an apartment in a block such that their ethnic quota is not reached. There are a few complications here: first-time applicants and low-income families usually receive priority numbers in the lottery scheme; moreover, the same estate may have several balloting rounds in order to ensure that all apartments are allocated by the time of completion. However, the focus of this work is on the welfare effects of using ethnic quotas rather than the specific intricacies of the HDB lottery mechanism. We must mention here the existing literature on the documentation of Singapore’s residential desegregation policies [16, 18, 38] and the empirical evaluation of their impact on various socioeconomic factors [41, 43]; to the best of our knowledge, ours is the first formal approach towards this problem.

2 PRELIMINARIES

We first describe a formal model for the housing allocation problem with ethnicity quotas. Throughout the paper, given $s \in \mathbb{N}$, we let $[s]$ denote the set $\{1, 2, \dots, s\}$.

Definition 2.1 (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 ,
- (iii) a utility $u(i, j) \in \mathbb{R}_+$ for each agent $i \in N$ and each item $j \in M$,
- (iv) a capacity $\lambda_{pq} \in \mathbb{N}$ for all $(p, q) \in [k] \times [l]$, indicating the upper bound on the number of agents of type N_p allowed in M_q .

We assume here that the inequality $\lambda_{pq} \leq |M_q|$ holds for all type-block pairs $(p, q) \in [k] \times [l]$ without loss of generality since it is not possible to assign more than $|M_q|$ agents of type N_p to items in block M_q , by definition. In the Singapore public allocation problem, tenant households are the agents and apartments are the items; types correspond to ethnic groups (Chinese, Malay, Indian/Others) and blocks to actual apartment blocks in a sales launch. In general, such partitions could be based on any criterion such as gender, profession, geographical location, or suchlike. For our theoretical analysis, we consider the idealized scenario where we have a central planner who has access to the utilities of each agent for all items, and determines an assignment that maximizes social welfare under type-block constraints.

An assignment of items to agents 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 ; a feasible solution is an assignment in which each item is allocated to at most one agent, and each agent receives at most one item, respecting the type-block capacities defined in (iv). We define the objective value as the utilitarian social welfare (or total utility), i.e. the sum of the utilities of all agents in an assignment $u(X) \triangleq \sum_{i \in N} \sum_{j \in M} x_{ij} u(i, j)$. Clearly, this optimization problem can be formulated as the following integer linear program:

$$\max \sum_{i \in N} \sum_{j \in M} x_{ij} u(i, j) \quad (1)$$

$$s.t. \quad \sum_{i \in N_p} \sum_{j \in M_q} x_{ij} \leq \lambda_{pq} \quad \forall p \in [k], \forall q \in [l] \quad (2)$$

$$\sum_{j \in M} x_{ij} \leq 1 \quad \forall i \in N \quad (3)$$

$$\sum_{i \in N} x_{ij} \leq 1 \quad \forall j \in M \quad (4)$$

$$x_{ij} \in \{0, 1\} \quad \forall i \in N, \forall j \in M \quad (5)$$

where constraints (3-5) jointly ensure that X is a matching of items to agents, and inequalities (2) embody our type-block constraints.

Finally, an instance of the decision version of ASSIGNTC consists of ingredients (i) to (iv) in Definition 2.1, as well as a positive value U : it is a ‘yes’-instance iff there exists a feasible assignment, satisfying constraints (2-5), whose objective value is at least U .

3 THE COMPLEXITY OF THE ASSIGNMENT PROBLEM WITH TYPE CONSTRAINTS

Our first main result is that the decision problem we introduce in Section 2 is NP-complete. We prove this by describing a polynomial-time reduction from the NP-complete Bounded Color Matching problem [22], defined as follows:

Definition 3.1 (BCMATCHING). An instance of the Bounded Color Matching (BCMATCHING) problem is given by (i) a bipartite graph $G = (A \cup B, E)$, where the set of edges E is partitioned into r subsets E_1, \dots, E_r representing r different edge colors, (ii) a capacity $w_t \in \mathbb{N}$ for each color $t \in [r]$, (iii) a profit $\pi_e \in \mathbb{Q}_+$ for each edge $e \in E$, and (iv) a positive integer P . It is a ‘yes’-instance iff there exists a matching (i.e. a collection of pairwise non-adjacent edges) $E' \subseteq E$ such that the sum of the profits of all edges in the matching is at least P , and there are at most w_t edges of color t in the matching, i.e. $\sum_{e \in E'} \pi_e \geq P$ and $|E' \cap E_t| \leq w_t$ for all $t \in [r]$.

THEOREM 3.2. *The ASSIGNTC problem is NP-complete.*

PROOF. That the problem is in NP is immediate: given an assignment, one can verify in poly-time that it satisfies the problem constraints and compute the social welfare. Given an instance $\langle G; \bar{w}; \bar{\pi}; P \rangle$ of BCMATCHING, we construct an instance of the ASSIGNTC problem as follows (see Example 3.3 for an illustration). Each edge $e \in E$ is an agent, whose type is its color. Items in our construction are partitioned into two blocks: M_1 and M_2 . The items in block M_1 correspond to the vertices in B : there is one item j_b for each node $b \in B$. For every $a \in A$, we add $\deg(a) - 1$ items $j_a^1, \dots, j_a^{\deg(a)-1}$ to M_2 , for a total of $|E| - |A|$ items. Thus, there is a total of $m = |B| + |E| - |A|$ items. Block M_1 accepts at most w_p agents of type N_p , whereas block M_2 has unlimited capacity; in other words, $\lambda_{p1} = w_p$ and $\lambda_{p2} = \min\{|N_p|, |M_2|\}$ for all $p \in [k]$. Given $e = (a, b)$, we define the utility function of agent e as follows:

$$u(e, j) = \begin{cases} \pi_e & \text{if } j = j_b, \\ \Phi & \text{if } j = j_a^s \text{ for some } s \in [\deg(a) - 1], \\ 0 & \text{otherwise.} \end{cases}$$

Here, Φ is an arbitrarily large constant, e.g., $\Phi = 1 + \sum_{e \in E} \pi_e$. Finally, let $U = P + \Phi(|E| - |A|)$. We begin by showing that if the BCMATCHING instance is a ‘yes’ instance, then so is our constructed ASSIGNTC instance. Let $E' \subseteq E$ be a matching whose value is at least P ; let us construct an assignment of items to agents via E' as follows. Observe some node $a \in A$; if $(a, b) \in E'$ then we assign the item $j_b \in M_1$ to the agent (a, b) ; the remaining $\deg(a) - 1$ agents of the form (a, b') , $b' \in B$, are arbitrarily assigned the items $j_a^1, \dots, j_a^{\deg(a)-1} \in M_2$. If E' contains no edges incident on a , then we arbitrarily choose $\deg(a) - 1$ edges incident on a and assign the corresponding agents to the items $j_a^1, \dots, j_a^{\deg(a)-1}$. We now show that it is a valid assignment satisfying the type-block constraints.

First, by construction, every agent (a, b) is assigned at most one item. Moreover, since E' is a matching, every item $j_b \in M_1$ is assigned to at most one agent of the form (a, b) ; hence, every item in M_2 is assigned to at most one agent by construction.

Let $E'_p = E_p \cap E'$ be the edges of color p in E' . Since matching E' satisfies the capacity constraints of the BCMATCHING instance, we have $|E'_p| \leq w_p$ for all $p \in [k]$; in particular, the number of items in M_1 assigned to agents of type p is no more than $w_p = \lambda_{p1}$. Thus, the type-block constraints for M_1 are satisfied. Moreover, the type-block constraints for M_2 are trivially satisfied. Hence our constructed assignment is valid and satisfies the type-block constraints.

Finally, we want to show that the social welfare exceeds U , the prescribed bound. Let us fix a node $a \in A$. By our construction, if the edge $e = (a, b)$ is in the matching E' , then agent e is assigned the item j_b for a utility of π_e . Thus the total utility of agents in E' equals $\sum_{e \in E'} \pi_e$, which is at least P by choice of E' . In addition, for every $a \in A$, there are exactly $\deg(a) - 1$ agents assigned to items in M_2 for a total utility of $\Phi(\deg(a) - 1)$. Summing over all $a \in A$, we have that the total utility derived by agents in $E \setminus E'$ is

$$\sum_{a \in A} \Phi(\deg(a) - 1) = \Phi \left(\sum_{a \in A} \deg(a) - \sum_{a \in A} 1 \right) = \Phi(|E| - |A|).$$

Putting it all together, we conclude that the total utility obtained by our assignment is at least $P + \Phi(|E| - |A|) = U$.

Next, we assume that our constructed ASSIGNTC instance is a ‘yes’ instance, and show that the original BCMATCHING instance must also be a ‘yes’ instance. Let X be a constrained assignment whose social welfare is at least $U = P + \Phi(|E| - |A|)$. Let E' be the set of edges corresponding to agents (a, b) assigned to items in M_1 ; we show that E' is a valid matching whose value is at least P . First, for any $b \in B$, X must assign the item j_b to at most one agent $e \in E'$. Next, since Φ is greater than the total utility obtainable from assigning all items in M_1 , it must be the case that X assigns all items $j_a^1, \dots, j_a^{\deg(a)-1}$ to $\deg(a) - 1$ agents of the form (a, b) , with $b \in B$, for every node $a \in A$; thus, there can be one edge in E' that is incident on a for every $a \in A$. Next, since X satisfies the type-block constraints, we know that for every $p \in [k]$, there are at most $\lambda_{p1} = w_p$ agents from E_p that are assigned items in M_1 ; thus, E' satisfies the capacity constraints. Finally, the utility extracted from the agents assigned to items in M_2 is exactly $\Phi(|E| - |A|)$; the total utility of the matching X is at least $U = P + \Phi(|E| - |A|)$, thus E' has a total profit of at least P in the BCMATCHING instance. \square

Example 3.3. In Figure 1, the graph $G = (A \cup B, E_1 \cup E_2)$, with $A = \{a_1, a_2\}$, $B = \{b_1, b_2, b_3\}$, $E_1 = \{(a_1, b_1), (a_2, b_2)\}$ and $E_2 = \{(a_1, b_2), (a_2, b_1), (a_2, b_3)\}$, is an instance of the BCMATCHING problem; edge labels are profits. The associated instance of the ASSIGNTC problem is defined by $N = N_1 \cup N_2$ and $M = M_1 \cup M_2$, where $N_1 = \{(a_1, b_1), (a_2, b_2)\}$, $N_2 = \{(a_1, b_2), (a_2, b_1), (a_2, b_3)\}$, $M_1 = \{j_{b_1}, j_{b_2}, j_{b_3}\}$ and $M_2 = \{j_{a_1}^1, j_{a_2}^1, j_{a_2}^2\}$; the utility of an agent for an item is equal to 0 if there is no edge between them, to Φ if the edge is dashed, and to the edge label otherwise.




Figure 1: A reduction from BCMATCHING to ASSIGNTC.

3.1 A Polynomial-Time Constant Factor Approximation Algorithm

Having established that the ASSIGNTC problem is computationally intractable in general, we next present an efficient constant-factor approximation algorithm: we construct an approximation-preserving reduction [35] – in fact, an S-reduction [17] – from this problem to the BCMATCHING problem (Definition 3.1), for which a polynomial-time approximation algorithm is known.

THEOREM 3.4. *There exists a poly-time $\frac{1}{2}$ -approximation algorithm for the ASSIGNTC problem.*

PROOF. Given an instance of the ASSIGNTC problem, we define a complete bipartite graph whose nodes correspond to the sets of agents N and items M , and give the edge joining agent-node i to item-node j a profit equal to the utility $u(i, j)$ for all $i \in N, j \in M$. We also give all edges joining agents of one type to items in one block the same color, so that there are kl colors indexed lexicographically by pairs $(p, q) \in [k] \times [l]$; let the capacity for color (p, q) be λ_{pq} . This produces, in $O(mn)$ time, an instance of BCMATCHING; the size of this instance is obviously polynomial in that of the original, and, by construction, there is a one-to-one correspondence between the sets of feasible solutions of the original and reduced instances with each corresponding pair having the same objective value (sum of edge-profits/utilities), so that the optimal values of the instances are also equal. We can now apply the polynomial-time $\frac{1}{2}$ -approximation algorithm introduced by Stamoulis [2014] for BCMATCHING on general weighted graphs. \square

Theorem 3.4 offers a $\frac{1}{2}$ -approximation to the ASSIGNTC problem; whether a better poly-time approximation algorithm exists is left for future work.

3.2 Uniformity Breeds Simplicity: Polynomial-Time Special Cases

Our results thus far make no assumptions on agent utilities; as we now show, the ASSIGNTC problem admits a poly-time algorithm under some assumptions on the utility model.

Definition 3.5 (Type-uniformity and Block-uniformity). A utility model u is called *type-uniform* if all agents of the same type have the same utility for an item, i.e. 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 $i \in N_p$. A utility model u is called *block-uniform* if all items in the same block offer the same utility to an agent; that is, 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 $j \in M_q$.

Type-uniformity assumes a strong correlation between agents’ type and utility; in the context of the HDB allocation problem, this implies that Singaporeans of the same ethnicity share the same preferences over apartments (perhaps due to cultural or socioeconomic factors). Cases that deal with uniform goods satisfy the block-uniformity assumption: e.g. students applying for spots in public schools or job applicants applying for multiple (identical) positions; in the HDB domain, block-uniformity captures purely location-based preferences, i.e. a tenant does not care which apartment she gets as long as it is in a specific block close to her workplace, family, or favorite public space.

THEOREM 3.6. *The ASSIGNTC problem can be solved in poly(n, m) time under either a type-uniform or a block-uniform utility model.*

We prove the result for a type-uniform utility model; the result for block-uniform utilities can be similarly derived. We propose a polynomial time algorithm based on the *Minimum-Cost Flow* problem which is known to be solvable in polynomial time. Recall that a flow network is a directed graph $G = (V, E)$ with a source node $s \in V$ and a sink node $t \in V$, where each arc $(a, b) \in E$

has a cost $\gamma(a, b) \in \mathbb{R}$ and a capacity $\psi(a, b) > 0$ representing the maximum amount that can flow on the arc; for convenience, we set $\gamma(a, b) = 0$ and $\psi(a, b) = 0$ for all $a, b \in V$ such that $(a, b) \notin E$. Let us denote by Γ and Ψ the matrices of costs and capacities respectively defined by $\Gamma = (\gamma(a, b))_{|V| \times |V|}$ and $\Psi = (\psi(a, b))_{|V| \times |V|}$. A flow in the network is a function $f : V \times V \rightarrow \mathbb{R}_+$ satisfying:

- (i) $f(a, b) \leq \psi(a, b)$ for all $a, b \in V$ (capacity constraints),
- (ii) $f(a, b) = -f(b, a)$ for all $a, b \in V$ (skew symmetry), and
- (iii) $\sum_{b \in V} f(a, b) = 0$ for all $a \in V \setminus \{s, t\}$ (flow conservation).

The value $v(f)$ of a flow f is defined by $v(f) = \sum_{a \in V} f(s, a) = \sum_{a \in V} f(a, t)$ and its cost $\gamma(f)$ is equal to $\sum_{(a, b) \in E} f(a, b)\gamma(a, b)$. The optimization problem can be formulated as follows. Given a value F , find a flow f that minimizes the cost $\gamma(f)$ subject to $v(f) = F$. This problem that takes as input the graph G , the matrices Γ and Ψ , and the value F , will be denoted by `MINCOSTFLOW` hereafter; given an instance $\langle G; \Gamma; \Psi; F \rangle$ of the `MINCOSTFLOW` problem, we let $\gamma(G, \Gamma, \Psi, F)$ be the cost of the optimal flow for that instance.

Given an instance \mathcal{I} of `ASSIGNTC`, we construct a flow network $G_{\mathcal{I}}(V, E)$ and matrices $\Gamma_{\mathcal{I}}$ and $\Psi_{\mathcal{I}}$ as follows (see Figure 2 for an illustration). The node set V is partitioned into layers: $V = \{s\} \cup A \cup B \cup C \cup \{t\}$. A is the *agent type layer*: there is one node $a_p \in A$ for each agent type $N_p, p \in [k]$. B is the *type-block layer*: it has a node $b_{pq} \in B$ for every type-block pair $(p, q) \in [k] \times [l]$. Finally, C is the *item layer*: there is one node $c_j \in C$ for all items $j \in M$. The arcs in E are as follows: for every a_p in A , there is an arc from s to a_p whose capacity $\psi(s, a_p)$ is $|N_p|$. Fixing $p \in [k]$, there is an arc from $a_p \in A$ to every $b_{pq} \in B$, where the capacity of (a_p, b_{pq}) is the capacity for type N_p in block M_q , i.e. $\psi(a_p, b_{pq}) = \lambda_{pq}$. Finally, given $q \in [l]$, there is an arc from b_{pq} to c_j iff $j \in M_q$; in this case, we have $\psi(b_{pq}, c_j) = 1$. The costs associated with arcs from B to C (i.e. arcs of the form (b_{pq}, c_j) where $j \in M_q$) are $-U_{pj}$; recall that U_{pj} is the utility that every agent of type N_p assigns to item j . All other arc costs are set to 0. We begin by proving a few technical lemmas on the above network.

Given a positive integer F , there exists an optimal flow that is integer-valued since $\langle G_{\mathcal{I}}; \Gamma_{\mathcal{I}}; \Psi_{\mathcal{I}}; F \rangle$ is integer-valued as well. Let f^* be an integer-valued optimal flow, taken over all possible values of F , that is:

$$f^* \in \operatorname{argmin}_{F \in [n]} \gamma(G_{\mathcal{I}}, \Gamma_{\mathcal{I}}, \Psi_{\mathcal{I}}, F) \quad (6)$$

Finding the flow f^* involves solving n instances of `MINCOSTFLOW` by definition; thus, one can find f^* in polynomial time. Given f^* as defined in (6), let $X^* = (x_{ij}^*)_{n \times m}$ be defined as follows: for every item $j \in M_q$, if $f^*(b_{pq}, c_j) = 1$ for some $p \in [k]$, then we choose an arbitrary unassigned agent $i \in N_p$ and set $x_{ij}^* = 1$.

LEMMA 3.7. X^* is a feasible solution of the `ASSIGNTC` instance \mathcal{I} .

PROOF. First, we assign at most one item to every agent by construction; next, let us show that each item $j \in M_q$ is assigned to at most one agent. Since f^* is a flow, we have $\sum_{p=1}^k f^*(b_{pq}, c_j) = f^*(c_j, t)$ due to flow conservation; note that the capacity of the arc (c_j, t) is 1, thus at most one arc (b_{pq}, c_j) has $f^*(b_{pq}, c_j) = 1$. Finally, since item j is assigned to an agent in N_p iff $f^*(b_{pq}, c_j) = 1$, we conclude that item j is assigned to at most one of the agents in N .




Figure 2: Network flow constructed for the proof of Theorem 3.6; in this case, we have 2 types and 2 blocks: $M_1 = \{1, 2\}$ and $M_2 = \{3, 4, 5\}$. Arc capacities are given in red. All arcs have a cost of 0, except those between $b_{pq} \in B$ and $c_j \in C$ whose cost equals $-U_{pj}$.

Next, let us prove that assignment X^* satisfies the type-block constraints; in other words, we need to show that:

$$\sum_{i \in N_p} \sum_{j \in M_q} x_{ij}^* \leq \lambda_{pq}, \forall p \in [k], \forall q \in [l] \quad (7)$$

Since f^* is a flow, we have $f^*(a_p, b_{pq}) = \sum_{j \in M_q} f^*(b_{pq}, c_j)$ for every pair $(p, q) \in [k] \times [l]$ due to flow conservation; moreover, we have $f^*(a_p, b_{pq}) \leq \psi(a_p, b_{pq}) = \lambda_{pq}$ by construction. Therefore, we necessarily have $\sum_{j \in M_q} f^*(b_{pq}, c_j) \leq \lambda_{pq}$ for all $p \in [k]$. Since an item $j \in M_q$ is matched with some agent $i \in N_p$ if and only if we have $f^*(b_{pq}, c_j) = 1$, we conclude that (7) indeed holds. \square

Now, let us establish a relation between the cost of f^* and the utility of the feasible assignment X^* .

LEMMA 3.8. The cost of the flow f^* satisfies $\gamma(f^*) = -u(X^*)$.

PROOF. By construction, the cost of f^* can only be induced by arcs from nodes in B to nodes in C , where the cost of all arcs of the form (b_{pq}, c_j) , with $j \in M_q$, is equal to $-U_{pj}$ (the negative of the uniform utility derived from item j by members of N_p). In other words, the cost of f^* can be written as follows:

$$\gamma(f^*) = - \sum_{p=1}^k \sum_{q=1}^l \sum_{j \in M_q} f^*(b_{pq}, c_j) U_{pj}$$

As previously argued, we have that $f^*(b_{pq}, c_j) \in \{0, 1\}$ for all arcs (b_{pq}, c_j) ; moreover, $f^*(b_{pq}, c_j) = 1$ iff item j is assigned to some agent in N_p . Therefore, we obtain:

$$\gamma(f^*) = - \sum_{p=1}^k \sum_{i \in N_p} \sum_{j \in M} x_{ij}^* U_{pj} = - \sum_{i \in N} \sum_{j \in M} x_{ij}^* u(i, j) = -u(X^*)$$

where the second equality holds since all agents in N_p have the same utility by assumption. \square

Finally, we show that for every feasible solution to the `ASSIGNTC` instance \mathcal{I} , there exists a flow with a matching cost.

LEMMA 3.9. *Let X be a feasible assignment for the ASSIGNTC instance \mathcal{I} ; there exists some feasible flow f such that $\gamma(f) = -u(X)$. Moreover, we have $v(f) = |\{i \in N : \sum_{j \in M} x_{ij} = 1\}|$.*

PROOF. Given a feasible assignment $X = (x_{ij})_{n \times m}$, we define $f : V \times V \rightarrow \mathbb{R}_+$ as follows:

$$\begin{cases} f(s, a_p) = \sum_{i \in N_p} \sum_{j \in M} x_{ij} & \forall a_p \in A \\ f(a_p, b_{pq}) = \sum_{i \in N_p} \sum_{j \in M_q} x_{ij} & \forall (a_p, b_{pq}) \in E \\ f(b_{pq}, c_j) = \sum_{i \in N_p} x_{ij} & \forall (b_{pq}, c_j) \in E \\ f(c_j, t) = \sum_{i \in N} x_{ij} & \forall c_j \in C \\ f(a, b) = -f(b, a) & \forall (a, b) \in E \\ f(a, b) = 0 & \forall (a, b) \notin E \end{cases}$$

The function f is indeed a flow: f trivially satisfies the skew symmetry condition by construction; next, we show that f satisfies flow conservation. For all $a_p \in A$, the incoming flow to node a_p from node s is $f(s, a_p) = \sum_{i \in N_p} \sum_{j \in M} x_{ij}$, and the outgoing flow to every b_{pq} is $\sum_{q=1}^l f(a_p, b_{pq}) = \sum_{i \in N_p} \sum_{j \in M} x_{ij}$ since M is partitioned into M_1, \dots, M_l ; hence flow is conserved. For a node $b_{pq} \in B$, the incoming flow equals $f(a_p, b_{pq}) = \sum_{i \in N_p} \sum_{j \in M_q} x_{ij}$ and an amount of $f(b_{pq}, c_j) = \sum_{i \in N_p} x_{ij}$ flows to every node c_j such that $j \in M_q$, thus flow is conserved. For a node $c_j \in C$ such that $j \in M_q$, its incoming flow equals $f(b_{pq}, c_j) = \sum_{i \in N_p} x_{ij}$ from every b_{pq} , for a total flow of $\sum_{p=1}^k \sum_{i \in N_p} x_{ij}$, which equals its outgoing flow to t . To conclude, f satisfies flow conservation.

Now let us prove that f satisfies the capacity constraints (i.e. $f(a, b) \leq \psi(a, b)$ for all arcs $(a, b) \in E$). For all $(s, a_p) \in E$, we have $f(s, a_p) = \sum_{i \in N_p} \sum_{j \in M} x_{ij} \leq |N_p| = \psi(s, a_p)$ since every agent $i \in N_p$ is matched with at most one item. For all $(a_p, b_{pq}) \in E$, we have $f(a_p, b_{pq}) = \sum_{i \in N_p} \sum_{j \in M_l} x_{ij} \leq \lambda_{pq} = \psi(a_p, b_{pq})$ since X satisfies the type-block constraints. For all arcs $(b_{pq}, c_j) \in E$, we have $f(b_{pq}, c_j) = \sum_{i \in N_p} x_{ij} \leq 1 = \psi(b_{pq}, c_j)$ since item j is matched with at most one of the agents in N_p . For all $(c_j, t) \in E$, we have $f(c_j, t) = \sum_{i \in N} x_{ij} \leq 1 = \psi(c_j, t)$ since item j is matched with at most one of the agents in N . Hence, f satisfies the capacity constraints and is a valid flow. Note that we have:

$$v(f) = \sum_{a \in V} f(s, a) = \sum_{p=1}^k f(s, a_p) = \sum_{p=1}^k \sum_{i \in N_p} \sum_{j \in M} x_{ij} = \sum_{i \in N} \sum_{j \in M} x_{ij}.$$

Then, since X is a feasible assignment of the ASSIGNTC instance \mathcal{I} , we conclude that we have $v(f) = |\{i \in N : \sum_{j \in M} x_{ij} = 1\}|$. We just need to prove that we have $\gamma(f) = -u(X)$, and we are done. By definition of the flow network, only arcs of the form (b_{pq}, c_j) contribute to the cost $\gamma(f)$ and we have $\gamma(b_{pq}, c_j) = -U_{pj}$; therefore, $\gamma(f) = -\sum_{(b_{pq}, c_j) \in E} f(b_{pq}, c_j) U_{pj}$. Since $f(b_{pq}, c_j) = \sum_{i \in N_p} x_{ij}$ (by definition of f) and $u(i, j) = U_{pj}$ for all agents $i \in N_p$ (by hypothesis), we finally obtain $\gamma(f) = -\sum_{j \in M} \sum_{p=1}^k \sum_{i \in N_p} x_{ij} u(i, j) = -\sum_{j \in M} \sum_{i \in N} x_{ij} u(i, j) = -u(X)$. \square

We are now ready to prove Theorem 3.6.

PROOF OF THEOREM 3.6. We begin by observing the flow f^* as defined in (6), and the assignment X^* derived from it. First, X^* is a feasible assignment of the ASSIGNTC instance \mathcal{I} (by Lemma 3.7) and we have $u(X^*) = -\gamma(f^*)$ (by Lemma 3.8). Then, for any valid assignment X of the ASSIGNTC instance \mathcal{I} , there exists a flow f

such that $\gamma(f) = -u(X)$; since $v(f) = |\{i \in N : \sum_{j \in M} x_{ij} = 1\}| \in [n]$, flow f is a feasible solution of the MINCOSTFLOW instance $\langle G_{\mathcal{I}}, \Gamma_{\mathcal{I}}, \Psi_{\mathcal{I}}, F \rangle$ for some $F \in [n]$. Therefore, we have:

$$u(X) = -\gamma(f) \leq -\gamma(G_{\mathcal{I}}, \Gamma_{\mathcal{I}}, \Psi_{\mathcal{I}}, v(f)) \leq -\gamma(f^*) = u(X^*)$$

Thus, X^* is an optimal solution of the ASSIGNTC instance \mathcal{I} ; since X^* can be computed in poly-time, this completes the proof. \square

4 THE PRICE OF DIVERSITY

We now turn to the allocative efficiency of the constrained assignment. As before, an instance of the ASSIGNTC problem is given by a set of n agents N partitioned into types N_1, \dots, N_k , a set of m items M partitioned into M_1, \dots, M_l , a list of capacity values $(\lambda_{pq})_{k \times l}$, and agent utilities for items given by $u = (u(i, j))_{n \times m}$. We denote the set of all assignments X of items to agents satisfying only the matching constraints (3-5) of Section 2 by \mathcal{X} , and that of all assignments additionally satisfying the type-block constraints (2) by \mathcal{X}_C ; the corresponding optimal social welfares for any given utility matrix $(u(i, j))_{n \times m}$ are:

$$OPT(u) \triangleq \max_{X \in \mathcal{X}} u(X); \quad OPT_C(u) \triangleq \max_{X \in \mathcal{X}_C} u(X).$$

Clearly, $OPT_C(u) \leq OPT(u)$ since $\mathcal{X}_C \subseteq \mathcal{X}$; we define the following natural measure of this welfare loss that lies in $[1, \infty]$:

Definition 4.1. For any instance of the ASSIGNTC problem, we define the *Price of Diversity* as follows, along the lines of Ahmed et al. [2017] and Bredereck et al. [2018]:

$$PoD(u) \triangleq OPT(u)/OPT_C(u).$$

The main result of this section is to establish an upper bound on $PoD(u)$ that is independent of the utility model. Denote the ratio of a type-block capacity to the size of the corresponding block by:

$$\alpha_{pq} \triangleq \lambda_{pq}/|M_q|.$$

THEOREM 4.2. *For any instance of ASSIGNTC, the tight upper bound (independent of utilities) on $PoD(u)$ is $1/\min_{(p,q) \in [k] \times [l]} \alpha_{pq}$.*

In general, the bound in Theorem 4.2 grows linearly in m (e.g. if the capacities λ_{pq} are fixed constants). However, type-block capacities are determined by a central planner in our model; a natural way of setting them is to fix the proportional capacities or quotas 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 $PoD(u)$ 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, potentially affecting the diversity objective adversely: it thus functions as a tunable tradeoff parameter between ethnic integration and worst-case welfare loss. In fact, in the Singapore allocation problem, the Ethnic Integration Policy fixes a universal percentage cap for each of the three ethnicities in all blocks; these percentages are set somewhat higher than the actual respective population proportions: the current block quotas α_{pq} are 0.87 for Chinese, 0.25 for Malays and 0.15 for Indian/Others [18]; plugging in these to the bound in Theorem 4.2, we have that the Singapore housing system has $PoD(u) \leq 6.67$. This bound makes no assumptions on agent utilities; in other words, it holds under *any utility model*.

The proof relies on the following lemma. Given an assignment $X \in \mathcal{X}$, let $u_p(X)$ denote the total utility of agents in N_p under X :

$$u_p(X) \triangleq \sum_{i \in N_p} \sum_{j \in M} x_{ij} u(i, j).$$

LEMMA 4.3. *For any instance of ASSIGNTC and any optimal unconstrained assignment $X^* \in \mathcal{X}$, we have:*

$$PoD(u) \leq u(X^*) / \left[\sum_{p \in [k]} u_p(X^*) \min_{q \in [l]} \alpha_{pq} \right].$$

PROOF SKETCH. Based on the optimal assignment X^* , we can construct an assignment $X \in \mathcal{X}_C$ satisfying the type-block constraints, by ‘revoking’ items in M_q from the agents with the smallest realized utilities in N_p whenever the type-block constraint is violated for (p, q) . We thus ensure that at least α_{pq} proportion of the utility remains under X for (p, q) . Summing over blocks, we obtain:

$$u_p(X) \geq u_p(X^*) \min_{q \in [l]} \alpha_{pq} \quad \forall p \in [k]. \quad (8)$$

Hence, since $X \in \mathcal{X}_C$, we have $PoD(u) \leq u(X^*)/u(X)$ where $u(X) = \sum_{p \in [k]} u_p(X)$. Combining this with inequality (8) gives is the required result. \square

We can now complete the proof of the theorem.

PROOF OF THEOREM 4.2. Since we have $\min_{(p,q) \in [k] \times [l]} \alpha_{pq} \leq \min_{q \in [l]} \alpha_{p'q}$ for all $p' \in [k]$, Lemma 4.3 immediately implies that $PoD(u) \leq 1/\min_{(p,q) \in [k] \times [l]} \alpha_{pq}$. Depending on the utility matrix u , this upper bound can be tight whenever $|N_{p_0}| \geq |M_{q_0}|$ for some type-block pair (p_0, q_0) in the set $\operatorname{argmin}_{(p,q) \in [k] \times [l]} \alpha_{pq}$. We identify a utility matrix for which the bound holds with equality:

$$u(i, j) = \begin{cases} 1 & \text{if } i \in N_{p_0} \text{ and } j \in M_{q_0}, \\ 0 & \text{otherwise.} \end{cases}$$

The optimal assignment without type-block constraints fully allocates the items in block M_{q_0} to agents in N_{p_0} for a total utility of $|M_{q_0}|$; furthermore, we know that 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 $\lambda_{p_0 q_0}$. Since $\lambda_{p_0 q_0} = \alpha_{p_0 q_0} \times |M_{q_0}|$, we have: $PoD(u) = |M_{q_0}|/\lambda_{p_0 q_0} = 1/\alpha_{p_0 q_0} = 1/\min_{(p,q) \in [k] \times [l]} \alpha_{pq}$. \square

4.1 The Impact of Ethnic Disparity

Theorem 4.2 offers a worst-case tight bound on the price of diversity, making no assumptions on agent utilities. However, its proof 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 obtain a better bound on the price of diversity if a less ‘disparate’ optimal assignment exists. To formalize this notion, we introduce a new parameter:

Definition 4.4. For an optimal unconstrained assignment $X^* \in \mathcal{X}$, denote by $\beta_p(X^*)$ 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 $\beta(X^*) \in (0, 1]$ can be computed in polynomial time and is fully independent of the type-block capacities. The closer $\beta(X^*)$ is to 1, the lower the disparity between average agents of different types under X^* .

THEOREM 4.5. *For any ASSIGNTC instance and any unconstrained optimal assignment $X^* \in \mathcal{X}$, we have:*

$$PoD(u) \leq 1 / \left[\beta(X^*) \sum_{p \in [k]} \frac{|N_p|}{n} \min_{q \in [l]} \alpha_{pq} \right].$$

PROOF. By definition of $\beta(X^*)$, we have, for every $p \in [k]$, $u_p(X^*) \geq \beta(X^*) \frac{|N_p|}{n} u(X^*)$. Substituting this in Lemma 4.3, we obtain the desired bound. \square

Let us now apply the result to the Singapore public housing domain; we use the ethnic proportions reported in the 2010 census report [19] to obtain $|N_1|/n = 0.741$ (Chinese), $|N_2|/n = 0.134$ (Malay), and $|N_3|/n = 0.125$ (Indian/Others). Using the same block quotas α_{pq} as before and assuming $\beta(X^*) = 1$, we have $PoD(u) \leq 1.43$. In general, by Theorems 4.2 and 4.5, if we plot the $PoD(u)$ against the disparity parameter $\beta(X^*)$, the point corresponding to any ASSIGNTC instance with above block quotas and ethnic proportions must lie in the shaded region of Figure 3.




Figure 3: PoD vs disparity parameter for the HDB problem.

5 EXPERIMENTAL ANALYSIS

We simulate instances of the ASSIGNTC problem using recent Singaporean demographic and housing allocation statistics. We compare the welfare of three assignment mechanisms: the optimal unconstrained mechanism, the optimal constrained mechanism, and the lottery-based mechanism used by HDB.

5.1 Data Collection

In order to create realistic instances of the ASSIGNTC problem within the Singaporean context, we collect data on the location and number of flats in recent HDB housing development projects advertised over the second and third quarters of 2017². Each of these developments corresponds to a block in our setup, for a total of $m = 1350$ flats partitioned into $l = 9$ blocks (a detailed map is given in Figure 4). We consider a pool of $n = 1350$ applicants whose ethnic composition follows the 2010 Singapore census report [19]: we have $|N_1| = 1000$ ($\approx 74.1\%$ Chinese), $|N_2| = 180$ ($\approx 13.4\%$ Malay) and $|N_3| = 170$ ($\approx 12.5\%$ Indian/Others). Finally, we use a uniform block capacity using the latest HDB block quotas [18]: for every block M_q , we have $\alpha_{1q} = 0.87$, $\alpha_{2q} = 0.25$ and $\alpha_{3q} = 0.15$.

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




Figure 4: Block locations and number of flats.

5.2 Utility Models

All parameters used to generate ASSIGNTC instances in our simulations are based on real data, except for agent utilities over apartments. Conducting large-scale surveys that elicit user preferences over apartments is beyond the scope of this work; thus, we base our agent utility models on simulated utilities. We examine two types of utility models: *distance-based* ($Dist(\sigma^2)$) and *ethnicity-based* ($Ethn(\sigma^2)$). In distance-based utilities, 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 , the utility agent i derives from apartments $j \in M_q$ is generated according to a normal distribution $\mathcal{N}(1/d(\vec{a}_i, loc(M_q)), \sigma^2)$; here, $loc(M_q)$ is the geographical location of block M_q , and $d(\cdot)$ is the Euclidean distance between \vec{a}_i and $loc(M_q)$. For the ethnicity-based utility model, we assume that all agents of the same ethnicity have the same preferred location (i.e. $\forall p \in [k], \forall i, i' \in N_p, \vec{a}_i = \vec{a}_{i'}$).

5.3 Evaluation

For both utility models, we vary σ^2 in $\{1, 5, 10\}$; the results reported in Figure 5 are on the average performance over 100 randomly generated instances. For each treatment, we report $PoD(u)$ the price of diversity (green), the theoretical upper bound on $PoD(u)$ as per Theorem 4.5 (blue), and the relative loss of the HDB lottery mechanism averaged over 50 runs (red).

We first observe that $Dist(\sigma^2)$ exhibits virtually no utility loss due to the imposition of type-block constraints. This is due to the fact that utilities generated according to the $Dist(\sigma^2)$ model are independent of ethnicities, resulting in a very low value for the inter-type disparity parameter (the blue column) — in fact, for any utility model where utilities are independent of ethnicities, the expected value of the disparity parameter should be 1. For utilities generated based on the $Ethn(\sigma^2)$ model, the disparity parameter is somewhat higher (utilities do strongly depend on ethnicity), resulting in a higher $PoD(u)$. Despite making no attempt to optimize social welfare under type-block constraints, the HDB lottery mechanism does surprisingly well, extracting approximately 85% of the optimal unconstrained welfare under the $Dist(\sigma^2)$ utility model, and at least 79% of the social welfare under the $Ethn(\sigma^2)$ model.

6 CONCLUSIONS AND FUTURE WORK

Our work constitutes a first step towards a better understanding of the effect that diversity constraints have on social welfare. We offer computational insights, providing a general hardness result, sufficient conditions for tractability, and a $\frac{1}{2}$ -approximation algorithm




Figure 5: Utility losses for $Dist(\sigma^2)$ (left) and $Ethn(\sigma^2)$ (right).

for ASSIGNTC. Our S-reduction essentially shows that ASSIGNTC is a special (still NP-complete) case of BCMATCHING and, although the question of a better approximation remains open, also implies some easy generalizations of our results. For example, there is a PTAS for BCMATCHING if one allows $(1 + \epsilon)$ -violations of the color constraints [23]; this immediately implies a PTAS for ASSIGNTC where one allows $(1 + \epsilon)$ -violations of the type-block constraints.

We derive two upper bounds on the price of diversity defined as the ratio of the optimal welfare achievable with and without type-block constraints: the first is in terms of block quotas only, independent of the utility model, hence under the planner's control; the second is parametrized in terms of inter-type disparity, which shows that when the disparity is low, the welfare loss is much closer to its ideal value of 1 than the first bound would suggest.

We analyze our model's behavior in simulation: the fundamental experimental framework is based on Singapore census and HDB sales data. Simulating agent utilities is still a major challenge: ideally, one would elicit applicants' utilities directly via large-scale national surveys. Our simulations tested two 'extreme' cases: one where there is no correlation between ethnicity and utility, and one where utility is artificially correlated to ethnicity. The truth is likely somewhere in between. Ethnic groups in Singapore most likely do have some correlation between their utility models; this can be due to socioeconomic factors (there is some correlation between ethnicity and socioeconomic status), the location of cultural or religious centers, or other unknown factors. Developing a more refined utility model is an interesting direction for future work.

Obviously, the HDB lottery mechanism cannot offer a better welfare than the allocation based on constrained optimization under known utilities; but, in our experiments, it does not perform significantly worse for the utility models we considered.

While our motivating problem is important in its own right, it is by no means an exclusive case where diversity constraints can have an effect on social welfare. Two immediate application domains are course allocation (with student capacity constraints), and public school allocation (with neighborhood location constraints).

Finally, our results describe an inevitable tradeoff between diversity and social welfare; however, we would like to strongly emphasize that this does not constitute a moral judgment on the authors' part. Economic considerations are certainly important, but they are by no means an exclusive nor a first order consideration. That said, understanding the impact of diversity constraints on social welfare is key if one is to justify their implementation.

ACKNOWLEDGMENTS

This research has been funded by an MOE Grant (no. R-252-000-625-133) and an NRF Research Fellowship (no. R-252-000-750-733).

REFERENCES

- [1] 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), 1954–1978.
- [2] Atila Abdulkadiroğlu and Tayfun Sönmez. 1998. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* 66, 3 (1998), 689–701.
- [3] Atila Abdulkadiroğlu and Tayfun Sönmez. 2003. School choice: A mechanism design approach. *American Economic Review* 93, 3 (2003), 729–747.
- [4] Faez Ahmed, John P Dickerson, and Mark Fuge. 2017. Diverse Weighted Bipartite b-Matching. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*. 35–41.
- [5] Kook Jin Ahn and Sudipto Guha. 2014. Near linear time schemes for uncapacitated and capacitated b -matching problems in nonbipartite graphs. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 239–258.
- [6] Elliot Anshelevich, Sanmay Das, and Yonatan Naamad. 2013. Anarchy, Stability, and Utopia: Creating Better Matchings. *Autonomous Agents and Multi-Agent Systems* 26, 1 (2013), 120–140.
- [7] Sophie Bade. 2016. *Random serial dictatorship: the one and only*. Working Paper. Royal Holloway College, University of London.
- [8] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. 2017. Finding Fair and Efficient Allocations. *arXiv preprint arXiv:1707.04731* (2017).
- [9] Siddharth Barman and Sanath Kumar Krishna Murthy. 2017. Approximation Algorithms for Maximin Fair Division. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*. 647–664.
- [10] Ulrich Bauer. 2004. *Assignment Problem with Constraints*. Master's thesis. ETH Zürich Department of Computer Science. Supervised by Andreas Weifl.
- [11] Anand Bhalgat, Deeparnab Chakrabarty, and Sanjeev Khanna. 2011. Social welfare in one-sided matching markets without money. In *Proceedings of the 14th International Workshop and 15th International Conference on Approximation, Randomization, and Combinatorial Optimization: Algorithms and Techniques (APPROX/RANDOM)*. 87–98.
- [12] Anna Bogomolnaia and Hervé Moulin. 2001. A new solution to the random assignment problem. *Journal of Economic Theory* 100, 2 (2001), 295–328.
- [13] Robert Bredereck, Piotr Faliszewski, Ayumi Igarashi, Martin Lackner, and Piotr Skowron. 2018. Multiwinner Elections with Diversity Constraints. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*. To appear.
- [14] Eric Budish and Estelle Cantillon. 2012. The multi-unit assignment problem: Theory and evidence from course allocation at Harvard. *American Economic Review* 102, 5 (2012), 2237–2271.
- [15] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. 2016. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*. 305–322.
- [16] Beng-Huat Chua. 1991. Race relations and public housing policy in Singapore. *Journal of Architectural and Planning Research* (1991), 343–354.
- [17] Pierluigi Crescenzi. 1997. A short guide to approximation preserving reductions. In *Proceedings of the 12th IEEE Conference on Computational Complexity (CCC)*. 262–273.
- [18] Yongheng Deng, Tien Foo Sing, and Chaoqun Ren. 2013. The story of Singapore's public housing: From a nation of home-seekers to a nation of homeowners. In *The Future of Public Housing: Ongoing Trends in the East and the West*, Jie Chen, Mark Stephens, and Yanyun Man (Eds.). Springer, Chapter 7, 103–121.
- [19] Department of Statistics, Singapore. 2010. Singapore 2010 Census: Key Indicators of the Resident Population. (2010).
- [20] Department of Statistics, Singapore. 2017. Singapore in Figures. (2017).
- [21] Clayton R. Featherstone. 2015. *Rank efficiency: Investigating a widespread ordinal welfare criterion*. Working Paper. The Wharton School, University of Pennsylvania.
- [22] Michael R. Garey and David S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co.
- [23] Fabrizio Grandoni, R. Ravi, and Mohit Singh. 2009. Iterative rounding for multi-objective optimization problems. In *Proceedings of the 17th Annual European Symposium on Algorithms (ESA)*. 95–106.
- [24] Housing and Development Board, Singapore. 2010. Policy changes to support an inclusive and cohesive home [Press release]. (05 March 2010). <http://www.nas.gov.sg/archivesonline/speeches/record-details/809e76bf-115d-11e3-83d5-0050568939ad>.
- [25] Housing and Development Board, Singapore. 2015. Costs and Fees. (07 October 2015). <http://www.hdb.gov.sg/cs/infoweb/residential/buying-a-flat/new/finance/costs-and-fees>.
- [26] Housing and Development Board, Singapore. 2017. Annual Report 2016/2017: Key Statistics. (2017). <http://www10.hdb.gov.sg/ebook/AR2017/key-statistics.html>.
- [27] Aanund Hylland and Richard Zeckhauser. 1979. The efficient allocation of individuals to positions. *Journal of Political Economy* 87, 2 (1979), 293–314.
- [28] Nicole Immorlica, Brendan Lucier, Joshua Mollner, and E. Glen Weyl. 2017. Raffles. (13 February 2017). Available at SSRN: <https://ssrn.com/abstract=2916337> or <http://dx.doi.org/10.2139/ssrn.2916337>.
- [29] Harold W. Kuhn. 1955. The Hungarian method for the assignment problem. *Naval Research Logistics (NRL)* 2, 1-2 (1955), 83–97.
- [30] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. 2016. When can the maximin share guarantee be guaranteed?. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, Vol. 16. 523–529.
- [31] Jérôme Lang and Piotr Krzysztof Skowron. 2016. Multi-Attribute Proportional Representation. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*. 530–536.
- [32] Jing Wu Lian, Nicholas Mattei, Renee Noble, and Toby Walsh. 2018. The Conference Paper Assignment Problem: Using Order Weighted Averages to Assign Indivisible Goods. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*. To appear.
- [33] László Lovász and Michael D. Plummer. 2009. *Matching theory*. Vol. 367. American Mathematical Society.
- [34] James Munkres. 1957. Algorithms for the assignment and transportation problems. *J. Soc. Indust. Appl. Math.* 5, 1 (1957), 32–38.
- [35] Pekka Orponen and Heikki Mannila. 1987. *On approximation preserving reductions: Complete problems and robust measures*. Technical Report C-1987-28. Department of Computer Science, University of Helsinki. 22 pages.
- [36] Parliament of Singapore. Parliament Debates: Official Report. 1989. Better racial mix in HDB housing estates. (16 February 1989). Vol. 52, cols. 650–668.
- [37] Parag A. Pathak and Tayfun Sönmez. 2013. School admissions reform in Chicago and England: Comparing mechanisms by their vulnerability to manipulation. *American Economic Review* 103, 1 (2013), 80–106.
- [38] Sock Yong Phang and Kyunghwan Kim. 2011. Singapore's Housing Policies: 1960–2013. *Frontiers in Development Policy: Innovative Development Case Studies* (2011), 123–153.
- [39] Ariel D. Procaccia and Junxing Wang. 2014. Fair enough: Guaranteeing approximate maximin shares. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*. 675–692.
- [40] Alvin E. Roth. 1984. The evolution of the labor market for medical interns and residents: a case study in game theory. *Journal of Political Economy* 92, 6 (1984), 991–1016.
- [41] Loo Lee Sim, Shi Ming Yu, and Sun Sheng Han. 2003. Public housing and ethnic integration in Singapore. *Habitat International* 27, 2 (2003), 293–307.
- [42] Georgios Stamoulis. 2014. Approximation algorithms for bounded color matchings via convex decompositions. In *Proceedings of the 39th International Symposium on Mathematical Foundations of Computer Science (MFCS)*. 625–636.
- [43] Maisy Wong. 2014. Estimating the distortionary effects of ethnic quotas in Singapore using housing transactions. *Journal of Public Economics* 115 (2014), 131–145.
- [44] Lin Zhou. 1990. On a conjecture by Gale about one-sided matching problems. *Journal of Economic Theory* 52, 1 (1990), 123–135.