Welfare Guarantees in Schelling Segregation

Martin Bullinger,[1] Warut Suksompong,[2] Alexandros A. Voudouris[3]

[1] Institut für Informatik, Technische Universität München
[2] School of Computing, National University of Singapore
[3] School of Computer Science and Electronic Engineering, University of Essex

Abstract

Schelling's model is an influential model that reveals how individual perceptions and incentives can lead to racial segregation. Inspired by a recent stream of work, we study welfare guarantees and complexity in this model with respect to several welfare measures. First, we show that while maximizing the social welfare is NP-hard, computing an assignment with approximately half of the maximum welfare can be done in polynomial time. We then consider Pareto optimality and introduce two new optimality notions, and establish mostly tight bounds on the worst-case welfare loss for assignments satisfying these notions. In addition, we show that for trees, it is possible to decide whether there exists an assignment that gives every agent a positive utility in polynomial time; moreover, when every node in the topology has degree at least 2, such an assignment always exists and can be found efficiently.

1 Introduction

Schelling's model was proposed half a century ago to illustrate how individual perceptions and incentives can lead to racial segregation, and has been used to study this phenomenon in residential metropolitan areas in particular (Schelling 1969, 1971). The model is rather simple to describe. There are a number of agents, each of whom belongs to one of two predetermined types and occupies a location; in his original work, Schelling assumed that the locations are cells of a rectangular board, which can be represented as a grid graph.[1] Every agent would like to occupy a node on the graph such that the fraction of other agents of the same type in the neighborhood of that node is at least a predefined tolerance threshold $\tau \in [0, 1]$. If this condition is not met for an agent, then the agent can relocate to a randomly chosen empty node on the grid. One of the most surprising findings of Schelling is that, starting from a random initial assignment of the agents to the nodes of the grid, the dynamics may converge to segregated assignments even when $\tau \approx 1/3$, contrasting the intuition that segregation should happen only when $\tau \geq 1/2$.

Throughout the years, hundreds of researchers in sociology and economics reconfirmed Schelling's observations and made similar ones for numerous variants of the model using computer simulations—see, e.g., (Clark and Fossett 2008). More recent work, mainly in computer science, performed rigorous analyses of such variants, some of which are quite close to the original model, and showed that the dynamics according to which the agents relocate converges to assignments in which the agents form large monochromatic regions (that is, subgraphs consisting only of agents of the same type); in addition, this line of work established bounds on the size of these regions. We refer to the papers (Pollicott and Weiss 2001; Young 2001; Zhang 2004; Pancs and Vriend 2007; Brandt et al. 2012; Barmpalias, Elwes, and Lewis-Pye 2014, 2015; Bhakta, Miracle, and Randall 2014; Immorlica et al. 2017) for results of this flavor.

While most of the literature on Schelling's model has focused on properties related to segregation between the two types, segregation itself is only one side of the story, especially when we allow different, possibly more complex location graphs. Given that the agents are willing to relocate to be close to other agents of the same type, another natural question is whether the resulting assignments satisfy some sort of *efficiency*. This has been considered in part by a recent array of papers (Chauhan, Lenzner, and Molitor 2018; Echzell et al. 2019; Elkind et al. 2019; Agarwal et al. 2020; Bilò et al. 2020; Chan, Irfan, and Than 2020; Kanellopoulos, Kyropoulou, and Voudouris 2020), which have studied Schelling's model from a game-theoretic perspective. In particular, instead of randomly relocating, the agents are assumed to be strategic and each of them aims to select a location that maximizes her *utility*, defined as the fraction of same-type agents in her neighborhood.

Besides questions related to the existence and computation of equilibria (i.e., assignments in which no agent has an incentive to relocate in order to increase her utility), the authors of some of the aforementioned papers have also studied the efficiency of assignments in terms of *social welfare*, defined as the total utility of the agents. For this objective, these authors have shown that computing assignments (not necessarily equilibria) maximizing the social welfare is NP-hard under specific assumptions about the graph and the behavior of the agents. Furthermore, they established several bounds on the worst-case ratio between the maximum social welfare (achieved by any possible assignment) and the social welfare of the best or worst equilbrium assignment, also

[1] Schelling also considered a variant in which the locations form a line graph.

known as the *price of stability* (Anshelevich et al. 2008) and the *price of anarchy* (Koutsoupias and Papadimitriou 1999), respectively. These ratios quantify the welfare that is lost due to the agents aiming to maximize their individual utilities rather than their collective welfare.

Inspired by this very recent stream of work, we study welfare guarantees and complexity in Schelling's model, not only with respect to the social welfare, but also to different notions of efficiency, such as Pareto optimality and natural variants of it.

1.1 Our Contribution

Our setting consists of n agents partitioned into two types, and a location graph known as the *topology*; agents of the same type are "friends", and agents of different types are "enemies". Each agent is assigned to a single node of the graph, and the utility of the agent is defined as the fraction of her friends among the agents in her neighborhood.

We start by considering the social welfare. We show that for any topology and any distribution of the agents into types, there always exists an assignment with social welfare at least $n/2 - 1$, and provide a polynomial-time algorithm for computing such an assignment. Since the social welfare never exceeds n, our algorithm produces an assignment with at least approximately half of the maximum social welfare. We complement this result by showing that maximizing the social welfare is NP-hard, even when the topology is a graph such that the number of nodes is equal to the number of agents. This improves upon previous hardness results of Elkind et al. (2019) and Agarwal et al. (2020) whose reductions use instances with "stubborn agents" (who are assigned to fixed nodes in advance and cannot move), and either a topology with the number of nodes larger than the number of agents, or at least three types of agents instead of just two. These results are presented in Section 3.

Even if an assignment does not maximize the social welfare, it does not mean that the assignment cannot be optimal in other senses. With this is mind, we next turn our attention to different notions of optimality. In particular, we consider the well-known notion of *Pareto optimality* (PO), according to which it should not be possible to improve the utility of an agent without decreasing that of another agent. We also introduce two variants of PO, called *utility-vector optimality* (UVO) and *group-welfare optimality* (GWO), which are particularly appropriate for Schelling's model but may be of interest in other settings as well. Informally, an assignment is UVO if we cannot improve the sorted utility vector of the agents, and GWO if it is not possible to increase the total utility of one type of agents without decreasing that of the other type. We prove several results on these three notions of optimality. First, while UVO and GWO imply PO by definition, we show that they are not implied by each other or by PO. Then, for each $X \in \{\mathrm{PO}, \mathrm{UVO}, \mathrm{GWO}\}$, we establish mostly tight bounds on the *price of* X, which is an analogue of the price of anarchy: the price of X is defined as the worst-case ratio between the maximum social welfare (among all assignments) and the minimum social welfare among all assignments satisfying X.[2] These results can be found in Section 4.

Finally, another important measure of efficiency is the number of agents who receive a positive utility in the assignment. Even though only requiring the utility to be nonzero seems minimal, there exist simple instances in which not all of the agents can be guaranteed to obtain a positive utility. We show that for trees, it is possible to decide in polynomial time whether there exists an assignment such that all agents receive a positive utility. In addition, we prove that it is always possible to guarantee a positive utility for at least half of the agents; moreover, when every node in the topology has degree at least 2, there exists an assignment in which all agents receive a positive utility, and such an assignment can be computed in polynomial time. These results are presented in Section 5.

1.2 Related Work

As already mentioned, Schelling's model and its variants have been studied extensively from many different perspectives in several disciplines. For an overview of early work on the model, we refer the reader to (Immorlica et al. 2017).

Most related to our present work are the papers (Elkind et al. 2019; Agarwal et al. 2020; Bilò et al. 2020; Kanellopoulos, Kyropoulou, and Voudouris 2020), which studied game-theoretic and complexity questions related to the social welfare in Schelling games. In particular, Elkind et al. (2019) considered *jump Schelling games* in which there are $k \geq 2$ types of agents, and the topology is a graph with more nodes than agents so that there are empty nodes to which unhappy agents can jump. They showed that equilibrium assignments do not always exist, proved that computing equilibrium assignments and assignments with social welfare close to n (the maximum possible) is NP-hard, and bounded the price of anarchy and stability for both general and restricted games.

Later on, Agarwal et al. (2020) considered the complement case of *swap Schelling games* in which the number of nodes in the topology is equal to the number of agents; since there are no empty nodes to which the agents can jump, the agents can increase their utility only by swapping positions pairwise. For this setting, the authors showed results similar to those of Elkind et al. (2019). Bilò et al. (2020) improved some of the price of anarchy bounds of Agarwal et al. (2020), and also studied a variation of the model in which the agents have a restricted view of the topology and can only swap with their neighbors. Finally, Kanellopoulos, Kyropoulou, and Voudouris (2020) investigated the price of anarchy and stability in jump Schelling games, but with a slightly different utility function according to which an agent considers herself as part of her set of neighbors.

The price of Pareto optimality was first considered by Elkind, Fanelli, and Flammini (2020), and implicitly stud-

[2] Note that an analogue of the price of *stability*, where we consider the worst-case ratio between the maximum social welfare and the *maximum* social welfare among assignments satisfying the optimality notion, is uninteresting: for all of the optimality notions we consider, this price is simply 1.

ied by Bullinger (2020), in the context of fractional hedonic games, which are closely related to Schelling games. Since Pareto optimality is a fundamental notion in various settings, its price has also been studied in the context of social distance games (Balliu, Flammini, and Olivetti 2017) and fair division (Bei et al. 2019). To the best of our knowledge, this is the first time that Pareto optimality is studied in Schelling's model.

2 Preliminaries

Let $N = \{1, \ldots, n\}$ be a set of $n \geq 2$ *agents*. The agents are partitioned into two different *types* (or *colors*), red and blue. Denote by r and b the number of red and blue agents, respectively; we have $r + b = n$. The distribution of agents into types is called *balanced* if $|r - b| \leq 1$. We say that two agents $i, j \in N$, $i \neq j$, are *friends* if i and j are of the same type; otherwise we say that they are *enemies*. For each $i \in N$, we denote the set of all friends of agent i by $F(i)$.

A *topology* is a simple connected undirected graph $G = (V, E)$, where $V = \{v_1, \ldots, v_t\}$. Each agent in N has to select a node of this graph so that there are no collisions. A tuple $I = (N, G)$ is called a *Schelling instance*. Given a set of agents N and a topology $G = (V, E)$ with $|V| \geq n$, an *assignment* is a vector $\mathbf{v} = (v(1), \ldots, v(n)) \in V^n$ such that $v(i) \neq v(j)$ for all $i, j \in N$ such that $i \neq j$; here, $v(i)$ is the node of the topology where agent i is positioned. A node $v \in V$ is *occupied* by agent i if $v = v(i)$. For a given assignment $\mathbf{v}$ and an agent $i \in N$, let $N_i(\mathbf{v}) = \{j \in N : \{v(i), v(j)\} \in E\}$ be the set of neighbors of agent i. Let $f_i(\mathbf{v}) = |N_i(\mathbf{v}) \cap F(i)|$ be the number of neighbors of i in $\mathbf{v}$ who are her friends. Similarly, let $e_i(\mathbf{v}) = |N_i(\mathbf{v})| - f_i(\mathbf{v})$ be the number of neighbors of i in $\mathbf{v}$ who are her enemies. Following prior work, we define the utility $u_i(\mathbf{v})$ of an agent i in $\mathbf{v}$ to be 0 if $N_i(\mathbf{v}) = 0$; otherwise, her utility is defined as the fraction of her friends among the agents in her neighborhood:

$$u_i(\mathbf{v}) = \frac{f_i(\mathbf{v})}{f_i(\mathbf{v}) + e_i(\mathbf{v})}.$$

The *social welfare* of an assignment $\mathbf{v}$ is defined as the total utility of all agents:

$$\mathrm{SW}(\mathbf{v}) = \sum_{i \in N} u_i(\mathbf{v}).$$

Let $\mathbf{v}^*(I)$ be an assignment that maximizes the social welfare for a given instance I; we refer to it as a *maximum-welfare* assignment. Note that for any assignment $\mathbf{v}$, we have $u_i(\mathbf{v}) \leq 1$, and so $\mathrm{SW}(\mathbf{v}^*) \leq n$. Denote by $\mathrm{SW}_R(\mathbf{v})$ and $\mathrm{SW}_B(\mathbf{v})$ the sum of the utilities of the red and blue agents, respectively; we have $\mathrm{SW}_R(\mathbf{v}) + \mathrm{SW}_B(\mathbf{v}) = \mathrm{SW}(\mathbf{v})$.

3 Social Welfare

The first question we address is whether a high social welfare can always be achieved in any Schelling instance. Even though it may seem that we can obtain high welfare simply by grouping the agents of each type together, given the possibly complex topology in combination with the distribution of agents into types, it is unclear how this idea can be executed in general or what guarantee it results in. Nevertheless, we show that high welfare is indeed always achievable. Moreover, we provide a tight lower bound on the maximum welfare for each number of agents.

For any positive integer n, define

$$g(n) = \begin{cases} \frac{n(n-2)}{2(n-1)} & \text{if } n \text{ is even;} \\ \frac{n-1}{2} & \text{if } n \text{ is odd.} \end{cases}$$

Note that $g(n) \geq n/2 - 1$ for all n. Our approach is to choose an assignment uniformly at random among all possible assignments. Equivalently, we place agents in the following iterative manner: for an arbitrary unoccupied node, assign a uniformly random agent who is unassigned thus far. We show that the expected welfare of the assignment resulting from this simple randomized algorithm is at least $g(n)$, which implies the existence of an assignment with this welfare guarantee.

Theorem 3.1. *For any Schelling instance with n agents, there exists an assignment with social welfare at least $g(n)$. Moreover, the bound $g(n)$ cannot be improved.*

Proof. First, note that we may assume that the number of agents is equal to the number of nodes by restricting our attention to an arbitrary connected subgraph of G with the desired size. For $v_i \in V$, let $N_{v_i} = \{v_j \in V \mid \{v_i, v_j\} \in E\}$ be the neighborhood of node v_i in G, and $n_{v_i} = |N_{v_i}|$ be its size.

Consider an assignment of the agents to the nodes of G chosen uniformly at random. Let W be a random variable denoting the social welfare of this assignment, U_i a random variable denoting the expected utility of the agent placed at node v_i, and X_i a binary random variable describing the color of this agent, where $X_i = 1$ if node v_i is occupied by a blue agent and $X_i = 0$ if it is occupied by a red agent. We have

$$\mathbb{E}[W] = \sum_{i=1}^{n} \mathbb{E}[U_i] = \sum_{i=1}^{n} (\Pr(X_i = 1) \cdot \mathbb{E}[U_i \mid X_i = 1] + \Pr(X_i = 0) \cdot \mathbb{E}[U_i \mid X_i = 0]).$$

In the first equality we use linearity of expectation, and in the second equality the law of total expectation.

Now, for a fixed $v_i \in V$, it holds that

$$\begin{aligned} \mathbb{E}[U_i \mid X_i = 1] &= \frac{1}{n_{v_i}} \sum_{v_j \in N_{v_i}} \mathbb{E}[X_j \mid X_i = 1] \\ &= \frac{1}{n_{v_i}} \sum_{v_j \in N_{v_i}} \Pr(v_j \text{ blue} \mid v_i \text{ blue}) \\ &= \frac{1}{n_{v_i}} \sum_{v_j \in N_{v_i}} \frac{b-1}{n-1} = \frac{b-1}{n-1}, \end{aligned}$$

where the first equality is again due to linearity of expectation. Similarly, we have $\mathbb{E}[U_i \mid X_i = 0] = \frac{r-1}{n-1}$. Hence,

$$\mathbb{E}[W] = \sum_{i=1}^{n} \left(\frac{b}{n} \cdot \frac{b-1}{n-1} + \frac{r}{n} \cdot \frac{r-1}{n-1} \right)$$

Algorithm 1 Assignment with high social welfare

Input: Schelling instance $I = (N, G = (V, E))$
Output: Assignment with social welfare at least $g(n)$

```
for i = 1, ..., n do
    if there is a unique assignment v consistent with
    X_1 = a_1, ..., X_{i-1} = a_{i-1} (up to permuting agents
    of the same color) then
        return v
    W_0 = E[W | X_1 = a_1, ..., X_{i-1} = a_{i-1}, X_i = 0]
    W_1 = E[W | X_1 = a_1, ..., X_{i-1} = a_{i-1}, X_i = 1]
    if W_1 ≥ W_0 then
        a_i = 1 /*assign a blue agent to v_i*/
    else
        a_i = 0 /*assign a red agent to v_i*/
return Assignment corresponding to (a_1, ..., a_n)
```

$$
\begin{aligned}
&= b \cdot \frac{b-1}{n-1} + r \cdot \frac{r-1}{n-1} \\
&= \frac{1}{n-1}(b(b-1) + (n-b)(n-b-1)) \\
&= \frac{1}{n-1}(n^2 - n + 2b(b-n)).
\end{aligned}
$$

Observe that the function $b(b-n)$ is decreasing in the range $b \in [0, n/2]$ and increasing in the range $b \in [n/2, n]$. This means that for even n, we have

$$
\begin{aligned}
\mathbb{E}[W] &\geq \frac{1}{n-1}\left(n^2 - n + 2 \cdot \frac{n}{2} \cdot \left(-\frac{n}{2}\right)\right) \\
&= \frac{n(n-2)}{2(n-1)} = g(n).
\end{aligned}
$$

For n odd, since b is an integer, it holds that

$$
\begin{aligned}
\mathbb{E}[W] &\geq \frac{1}{n-1}\left(n^2 - n + 2 \cdot \frac{n-1}{2} \cdot \left(-\frac{n+1}{2}\right)\right) \\
&= \frac{n-1}{2} = g(n),
\end{aligned}
$$

implying that $\mathbb{E}[W] \geq g(n)$ in both cases.

Finally, it can be verified that when G is a complete graph and the distribution of agents into types is balanced, every assignment has social welfare exactly $g(n)$. $\square$

Next, we derandomize the algorithm in Theorem 3.1 to produce an efficient deterministic algorithm that computes an assignment with welfare at least $g(n)$. The pseudocode of the algorithm, which shares the notation of Theorem 3.1, can be found in Algorithm 1. The idea is that when we choose an agent to be assigned to an unassigned node, we pick a type such that the expected welfare is maximized, where the expectation is taken with respect to the uniform distribution of the remaining agents to the remaining nodes.

Theorem 3.2. *Algorithm 1 returns an assignment with social welfare at least $g(n)$ in polynomial time.*

Proof. First, we prove that the welfare of the returned assignment is at least $g(n)$. For $i = 0, \ldots, n$, denote by A_i the event $X_1 = a_1 \wedge X_2 = a_2 \wedge \cdots \wedge X_i = a_i$. In particular, A_0 is the entire sample space. We will show by induction that for each i, $\mathbb{E}[W \mid A_i] \geq \mathbb{E}[W]$. The base case $i = 0$ holds trivially. For $i \in \{1, \ldots, n\}$, if there is a unique assignment consistent with $X_1 = a_1 \wedge \cdots \wedge X_{i-1} = a_{i-1}$, then the social welfare of the returned assignment is $\mathbb{E}[W \mid A_{i-1}] \geq \mathbb{E}[W] \geq g(n)$, where the first inequality follows from the induction hypothesis and the second inequality from Theorem 3.1. Otherwise, we have

$$
\begin{aligned}
\mathbb{E}[W] &\leq \mathbb{E}[W \mid A_{i-1}] \\
&= \Pr(X_i = 0 \mid A_{i-1}) \cdot \mathbb{E}[W \mid A_{i-1} \wedge X_i = 0] \\
&\quad + \Pr(X_i = 1 \mid A_{i-1}) \cdot \mathbb{E}[W \mid A_{i-1} \wedge X_i = 1] \\
&\leq \Pr(X_i = 0 \mid A_{i-1}) \cdot \mathbb{E}[W \mid A_i] \\
&\quad + \Pr(X_i = 1 \mid A_{i-1}) \cdot \mathbb{E}[W \mid A_i] \\
&= \mathbb{E}[W \mid A_i],
\end{aligned}
$$

where we use the law of total expectation for the first equality and the choice of a_i in the algorithm for the second inequality. This completes the induction. Hence, if the algorithm terminates in the jth iteration, the welfare of the returned assignment is $\mathbb{E}[W \mid A_j] \geq \mathbb{E}[W] \geq g(n)$.

We next show that the algorithm can be implemented in polynomial time. To this end, it suffices to show that the quantities W_0 and W_1 can be computed efficiently for each fixed $i \in \{1, \ldots, n\}$. If there is only one type of agents left after having assigned the first i agents, this is straightforward, so assume that both types of agents still remain. By linearity of expectation, for each $x \in \{0, 1\}$,

$$
\mathbb{E}[W \mid A_{i-1} \wedge X_i = x] = \sum_{j=1}^{n} \mathbb{E}[U_j \mid A_{i-1} \wedge X_i = x].
$$

By the law of total expectation,

$$
\begin{aligned}
\mathbb{E}[U_j \mid &A_{i-1} \wedge X_i = x] \\
&= \Pr(X_j = 0 \mid A_{i-1} \wedge X_i = x) \\
&\qquad \cdot \mathbb{E}[U_j \mid A_{i-1} \wedge X_i = x \wedge X_j = 0] \\
&+ \Pr(X_j = 1 \mid A_{i-1} \wedge X_i = x) \\
&\qquad \cdot \mathbb{E}[U_j \mid A_{i-1} \wedge X_i = x \wedge X_j = 1],
\end{aligned}
$$

where a probability can be 0 if v_j has already been assigned an agent (i.e., if $j \leq i$). When $j > i$, we have

$$
\Pr(X_j = 1 \mid A_{i-1} \wedge X_i = x) = \frac{b - \sum_{k=1}^{i-1} a_i - x}{n-i}.
$$

Also, by linearity of expectation,

$$
\begin{aligned}
&\mathbb{E}[U_j \mid A_{i-1} \wedge X_i = x \wedge X_j = 1] \\
&\quad = \frac{1}{n_{v_j}} \sum_{v_k \in N_{v_j}} \mathbb{E}[X_k \mid A_{i-1} \wedge X_i = x \wedge X_j = 1].
\end{aligned}
$$

Finally,

$$
\begin{aligned}
&\mathbb{E}[X_k \mid A_{i-1} \wedge X_i = x \wedge X_j = 1] \\
&\quad = \begin{cases} a_k & \text{if } k \leq i-1; \\ x & \text{if } k = i; \\ \frac{b - \sum_{\ell=1}^{i-1} a_\ell - x - 1}{n-i-1} & \text{if } k > i. \end{cases}
\end{aligned}
$$

The computations for $X_j = 0$ as well as for $j \leq i$ can be done similarly. $\square$

Since the social welfare of any assignment is at most n, Algorithm 1 always produces an assignment with at least roughly half of the optimal welfare. This raises the question of whether it is possible to compute a maximum-welfare assignment for any given instance in polynomial time. Unfortunately, Elkind et al. (2019, Thm. 4.2) proved that maximizing social welfare is NP-hard. However, their proof relies on the existence of a "stubborn agent", who is assigned to a fixed node in advance and cannot move, and uses a topology with more nodes than agents.[3] We show that the hardness remains even when both of these assumptions are removed and the topology is a regular graph.

Theorem 3.3. *It is NP-complete to decide whether there exists an assignment with social welfare at least s, given a Schelling instance and a rational number s, even for the class of instances where the number of agents is equal to the number of nodes and the topology is a regular graph.*

Due to space constraints, the proof of Theorem 3.3 (and all other missing proofs) can be found in the full version of our paper (Bullinger, Suksompong, and Voudouris 2020).

4 Optimality Notions

Even when an assignment does not have the maximum social welfare, there can still be other ways in which it is "optimal". In this section, we consider some optimality notions and quantify them in relation to social welfare. A classical optimality notion is Pareto optimality.

Definition 4.1. An assignment $\mathbf{v}$ is said to be *Pareto dominated* by assignment $\mathbf{v}'$ if $u_i(\mathbf{v}) \leq u_i(\mathbf{v}')$ for all $i \in N$, with the inequality being strict for at least one i. An assignment $\mathbf{v}$ is *Pareto optimal (PO)* if it is not Pareto dominated by any other assignment.

Given two vectors $\mathbf{w}_1$ and $\mathbf{w}_2$ of the same length k, we say that $\mathbf{w}_1$ *weakly dominates* $\mathbf{w}_2$ if for each $i \in \{1, \dots, k\}$, the ith element of $\mathbf{w}_1$ is at least that of $\mathbf{w}_2$. We say that $\mathbf{w}_1$ *strictly dominates* $\mathbf{w}_2$ if at least one of the inequalities is strict.

For an assignment $\mathbf{v}$, denote by $\mathbf{u}(\mathbf{v})$ the vector of length n consisting of the agents' utilities $u_i(\mathbf{v})$, sorted in non-decreasing order. Similarly, denote by $\mathbf{u}_R(\mathbf{v})$ and $\mathbf{u}_B(\mathbf{v})$ the corresponding vectors of length r and b for the red and blue agents, respectively. Note that an assignment $\mathbf{v}$ is Pareto optimal if and only if there is no other assignment $\mathbf{v}'$ such that $\mathbf{u}_X(\mathbf{v}')$ weakly dominates $\mathbf{u}_X(\mathbf{v})$ for $X \in \{R, B\}$ and at least one of the dominations is strict.

Motivated by this observation, we define two new optimality notions appropriate for Schelling instances.

Definition 4.2. An assignment $\mathbf{v}$ is said to be

- *group-welfare dominated* by assignment $\mathbf{v}'$ if $\mathrm{SW}_X(\mathbf{v}') \geq \mathrm{SW}_X(\mathbf{v})$ for $X \in \{R, B\}$ and at least one of the inequalities is strict;
- *utility-vector dominated* by assignment $\mathbf{v}'$ if $\mathbf{u}(\mathbf{v}')$ strictly dominates $\mathbf{u}(\mathbf{v})$.

[3]Agarwal et al. (2020) showed that hardness holds when the numbers of agents and nodes are equal, but still required stubborn agents and moreover assumed at least three types of agents.

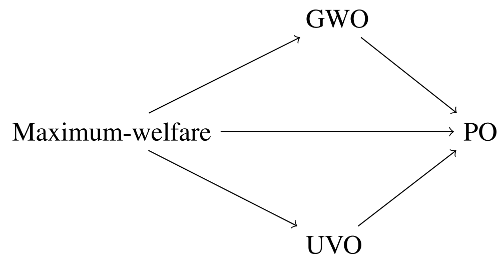

Figure 1: Implication relations among optimality notions

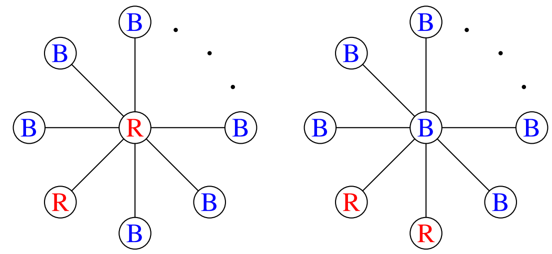

Figure 2: Example showing that GWO does not imply UVO.

An assignment $\mathbf{v}$ is *group-welfare optimal (GWO)* (resp., *utility-vector optimal (UVO)*) if it is not group-welfare dominated (resp., utility-vector dominated) by any other assignment.

The implication relations in Figure 1 follow immediately from the definitions; in particular, both of the new notions lie between welfare maximality and Pareto optimality. We claim that no other implications exist between these notions. To establish this claim, it suffices to show that GWO and UVO do not imply each other.

Proposition 4.3. *GWO does not imply UVO.*

Proof. Assume that the topology is a star as in Figure 2, and there are two red and $n-2$ blue agents, where $n \geq 5$. The left assignment $\mathbf{v}$ is GWO, since putting a blue agent at the center as in the right assignment $\mathbf{v}'$ leaves both red agents with utility 0. However, $\mathbf{v}$ is not UVO, as

$$\mathbf{u}(\mathbf{v}) = (1, 1/(n-1), 0, \dots, 0)$$

is strictly dominated by

$$\mathbf{u}(\mathbf{v}') = (1, \dots, 1, (n-3)/(n-1), 0, 0). \qquad \square$$

Proposition 4.4. *UVO does not imply GWO.*

Proof. Let n be a multiple of 4. Suppose that the topology is a complete bipartite graph with $n/2$ nodes on each side, and there are $n/2$ red and $n/2$ blue agents (Figure 3). The left assignment $\mathbf{v}$, which assigns one red agent to the left side and one blue agent to the right side, is UVO. Indeed, the red agent assigned to the left side receives utility $(n/2-1)/(n/2)$, and any assignment in which an agent receives equal or higher utility must have the same sorted utility vector as $\mathbf{v}$. We have

$$\mathrm{SW}(\mathbf{v}) = 2 \cdot \frac{n/2-1}{n/2} + 2(n/2-1) \cdot \frac{1}{n/2} = 4 - \frac{8}{n},$$

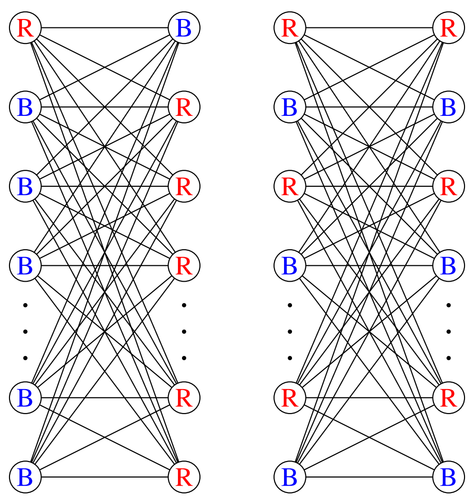

Figure 3: Example showing that UVO does not imply GWO. The topology is a complete bipartite graph.

with each group receiving half of the welfare, i.e., $2 - 4/n$. On the other hand, in the right assignment $\mathbf{v}'$, which assigns half of the agents of each color to each side, every agent receives utility $1/2$. Hence $\text{SW}(\mathbf{v}') = n/2$, and each group receives a total utility of $n/4$. Hence, when $n \geq 8$, $\mathbf{v}$ is UVO but not GWO. □

In order to quantify the welfare guarantee that each optimality notion provides, we define the price of a notion as follows.

Definition 4.5. Given a property P of assignments and a Schelling instance, the *price of* P for that instance is defined as the ratio between the maximum social welfare and the minimum social welfare of an assignment satisfying P:

$$\textit{Price of } P \textit{ for instance } I = \frac{\text{SW}(\mathbf{v}^*(I))}{\min_{\mathbf{v} \in P(I)} \text{SW}(\mathbf{v})},$$

where $P(I)$ is the set of all assignments satisfying P in instance I.[4]

The *price of* P for a class of instances is then defined as the supremum price of P over all instances in that class.

For $P \in \{PO, GWO, UVO\}$, we have $\mathbf{v}^*(I) \in P(I)$, so the price of P is always well-defined and at least 1. Note also that $\max_{\mathbf{v} \in P(I)} \text{SW}(\mathbf{v}) = \text{SW}(\mathbf{v}^*(I))$.

In Figure 2, the left assignment is GWO and PO and has social welfare $n/(n-1)$, whereas the maximum-welfare assignment on the right has social welfare $n(n-3)/(n-1)$. We therefore have the following bound (for $n \leq 4$, the bound holds trivially).

Proposition 4.6. *For each n, the price of GWO and the price of PO are at least $n - 3$.*

The following result shows that the welfare of a UVO assignment can also be a linear factor away from the maximum welfare, but not more.

[4]We interpret the ratio $\frac{0}{0}$ in this context to be equal to 1.

Theorem 4.7. *The price of UVO is $\Theta(n)$.*

Proof. Lower bound: Consider the topology in Figure 3. As in the proof of Proposition 4.4, the left assignment $\mathbf{v}$ is UVO and has social welfare $4 - 8/n$. On the other hand, the right assignment $\mathbf{v}'$ has social welfare $n/2$, meaning that the ratio $\text{SW}(\mathbf{v}')/\text{SW}(\mathbf{v})$ is greater than $n/8$.

Upper bound: We claim that if $n \geq 3$, any UVO assignment has social welfare at least[5] $1/2$; since the maximum social welfare is at most n, this yields the desired bound. Assume first that the number of agents is equal to the number of nodes. Let $\mathbf{v}$ be a UVO assignment. If there is a red agent and a blue agent both receiving utility 0, then since no node is empty and $n \geq 3$, swapping them yields an improvement with respect to the utility vector. So we may assume that all agents of one type, say blue, receive a positive utility. If at least $n/2$ agents receive a positive utility, then $\text{SW}(\mathbf{v}) \geq n/(2n-2) \geq 1/2$. Assume therefore that more than $n/2$ agents receive utility 0; these agents must all be red. Swap b of these red agents with utility 0 with all b blue agents to obtain an assignment $\mathbf{v}'$. Each of these b red agents receives utility in $\mathbf{v}'$ at least the utility in $\mathbf{v}$ of the blue agent with whom it was swapped, while all blue agents receive utility 0 in $\mathbf{v}'$. Every other (red) agent is not worse off, and at least one of them is better off (in particular, one who receives utility 0 in $\mathbf{v}$, which must exist since $n/2 > b$). Hence $\mathbf{v}$ is utility-vector dominated by $\mathbf{v}'$, a contradiction.

Now, assume that the number of agents is less than the number of nodes. Since $n \geq 3$, any UVO assignment $\mathbf{v}$ must have $\text{SW}(\mathbf{v}) > 0$, so there exists a connected component of $\mathbf{v}$ with a positive social welfare. Let n' be the size of this component. If $n' = 2$, then $\text{SW}(\mathbf{v}) \geq 2$. Else, the assignment restricted to this component is also UVO, and by our earlier arguments has social welfare at least $1/2$. □

Next, we show that the price of GWO is also linear. We start by establishing a lower bound on the social welfare of GWO assignments.

Theorem 4.8. *Any GWO assignment has social welfare at least $n/(n-1)$ for $n \geq 4$, and 1 for $n = 3$. Moreover, these bounds cannot be improved.*

Since the social welfare never exceeds n, Proposition 4.6 and Theorem 4.8 imply that the price of GWO is $\Theta(n)$.

We now turn to Pareto optimality, for which we prove a weaker lower bound on the social welfare.

Theorem 4.9. *When $n \geq 3$, any PO assignment has social welfare at least $1/\sqrt{n}$.*

Combined with Proposition 4.6, Theorem 4.9 implies that when $n \geq 3$, the price of PO is at least $n - 3$ and at most $n\sqrt{n}$.

We conjecture that the welfare guarantee in Theorem 4.9 can be improved to $n/(n-1)$ for $n \geq 4$, which would be tight due to the left assignment in Figure 2. In our full version, we confirm this conjecture when the topology is a tree; this also implies that the price of PO is $\Theta(n)$ in this special

[5]In the full version of our paper, we improve this bound to 1 via a longer proof (Bullinger, Suksompong, and Voudouris 2020).

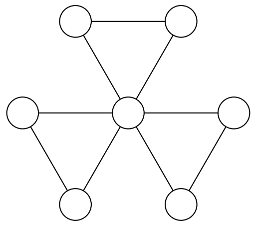

Figure 4: Example showing that Theorem 5.3 does not hold when the number of nodes is greater than the number of agents. There are three red and three blue agents. No matter how the agents are placed, at least one of them will receive utility 0.

case. In addition, we prove a linear bound on the price of PO when the fraction of agents of each type is at least some constant.

5 Number of Positive Agents

In this section, we consider the problem of maximizing the number of agents receiving a positive utility, who we refer to as *positive agents*. Notice that it is not always possible to make every agent positive—for example, in a star, every agent whose type is different from the center agent receives zero utility. We begin by showing that for trees, deciding whether it is possible to make every agent positive can be done efficiently. Our algorithm is based on dynamic programming and shares some similarities with the algorithm of Elkind et al. (2019) for deciding whether an equilibrium exists on a tree.

Theorem 5.1. *There is a polynomial-time algorithm that decides whether there exists an assignment in which every agent receives a positive utility when the topology is a tree.*

Observe that for any topology, an assignment in which at least half of the agents are positive is guaranteed to exist and can be easily found by using depth-first search for the majority type.

Proposition 5.2. *For any $n \geq 3$, there exists a polynomial-time algorithm that computes an assignment in which at least $\lceil n/2 \rceil$ agents receive a positive utility.*

The bound $\lceil n/2 \rceil$ is tight when the topology is a star and there are $\lceil n/2 \rceil$ red and $\lfloor n/2 \rfloor$ blue agents.

Next, we show that when every node has degree at least 2 and the number of agents is equal to the number of nodes, it is possible to give every agent a positive utility. Note that the latter condition is also necessary—for the topology given in Figure 4, if there are three red and three blue agents (so one node is left unoccupied), it is easy to see that no assignment makes every agent positive.

Theorem 5.3. *Suppose that every node in the topology has degree at least 2, the number of agents is equal to the number of nodes, and there are at least two agents of each type. Then there exists an assignment such that every agent receives a positive utility.*

Proof. Consider an arbitrary assignment $\mathbf{v}$. If every agent is already positive, we are done, so assume that there is an agent i with utility 0. Without loss of generality, i is a blue agent. Among all paths from i to another blue agent, consider one with maximum length—suppose that the path goes to agent j. Since there are at least two blue agents, such a path must exist; moreover, since i has utility 0, the path contains at least one red agent. Let k be the last red agent on the path before reaching j. Swap i and k.

We claim that in the resulting assignment $\mathbf{v}'$, the number of positive agents increases by at least 1; by applying such swaps repeatedly, we will reach an assignment in which all agents are positive. To establish the claim, it suffices to show that i, k, as well as any agent adjacent to either of them are positive in $\mathbf{v}'$. Since i has utility 0 in $\mathbf{v}$, she has at least two red neighbors in $\mathbf{v}$, so k is positive in $\mathbf{v}'$. Moreover, i is adjacent to j in $\mathbf{v}'$ and therefore becomes positive. Any other red agent on the path remains positive, and all agents adjacent to k in $\mathbf{v}'$ are red (besides possibly i, if k is the only red agent on the path) and are therefore positive. Finally, consider any red agent ℓ adjacent to i in $\mathbf{v}'$ not lying on the path. Since every node has degree at least 2, agent ℓ must have a neighbor $m \neq i$ (possibly $m = j$). If ℓ is adjacent to k, then ℓ is positive since k is red. Else, if m is a blue agent, we obtain a longer path from i to m in $\mathbf{v}$ than the original longest path, a contradiction. Hence m must be red, and ℓ is positive in $\mathbf{v}'$, proving the claim. $\square$

Since the longest path problem is known to be NP-hard, the proof of Theorem 5.3 does not give rise to a polynomial-time algorithm for computing a desired assignment. In the full version of our paper, we present an inductive approach that is more involved but leads to an efficient algorithm (Bullinger, Suksompong, and Voudouris 2020).

6 Conclusion

In this paper, we have studied questions regarding welfare guarantees and complexity in Schelling segregation. Our findings are mostly positive: An assignment with high social welfare always exists, and the welfare of assignments satisfying most optimality notions are at most a linear factor away from the maximum social welfare. Furthermore, even though an assignment yielding a positive utility to every agent may not exist, the existence can be guaranteed when every node in the topology has degree at least 2, a realistic assumption in well-connected metropolitan areas.

Several interesting directions remain from our work. On the technical side, it would be useful to close the gap on the price of Pareto optimality, which we conjecture to be $\Theta(n)$, and characterize the topologies for which an assignment such that every agent receives a positive utility always exists. Another open question is whether we can obtain a better approximation to social welfare in polynomial time. From a more conceptual perspective, one could try to extend our results to a model with more than two types of agents or more complex friendship relations (Elkind et al. 2019) or modified utility functions (Kanellopoulos, Kyropoulou, and Voudouris 2020). Studying our new optimality notions from Section 4 in related settings such as hedonic games, especially when agents are partitioned into types, may lead to intriguing discoveries as well.

Acknowledgments

This work was partially supported by the Deutsche Forschungsgemeinschaft under grant BR 2312/12-1, by the European Research Council (ERC) under grant number 639945 (ACCORD), and by an NUS Start-up Grant. We would like to thank Jiarui Gan, Pascal Lenzner, and Louise Molitor for interesting discussions, and the anonymous reviewers for helpful comments.

References

Agarwal, A.; Elkind, E.; Gan, J.; and Voudouris, A. A. 2020. Swap stability in Schelling games on graphs. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 1758–1765.

Anshelevich, E.; Dasgupta, A.; Kleinberg, J. M.; Tardos, É.; Wexler, T.; and Roughgarden, T. 2008. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing* 38(4): 1602–1623.

Balliu, A.; Flammini, M.; and Olivetti, D. 2017. On Pareto optimality in social distance games. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 349–355.

Barmpalias, G.; Elwes, R.; and Lewis-Pye, A. 2014. Digital morphogenesis via Schelling segregation. In *Proceedings of the 55th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, 156–165.

Barmpalias, G.; Elwes, R.; and Lewis-Pye, A. 2015. Unperturbed Schelling segregation in two or three dimensions. *Journal of Statistical Physics* 164: 1460–1487.

Bei, X.; Lu, X.; Manurangsi, P.; and Suksompong, W. 2019. The price of fairness for indivisible goods. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 81–87.

Bhakta, P.; Miracle, S.; and Randall, D. 2014. Clustering and mixing times for segregation models on $\mathbb{Z}^2$. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 327–340.

Bilò, D.; Bilò, V.; Lenzner, P.; and Molitor, L. 2020. Topological influence and locality in swap Schelling games. In *Proceedings of the 45th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, 15:1–15:15.

Brandt, C.; Immorlica, N.; Kamath, G.; and Kleinberg, R. 2012. An analysis of one-dimensional Schelling segregation. In *Proceedings of the 44th Symposium on Theory of Computing Conference (STOC)*, 789–804.

Bullinger, M. 2020. Pareto-optimality in cardinal hedonic games. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 213–221.

Bullinger, M.; Suksompong, W.; and Voudouris, A. A. 2020. Welfare guarantees in Schelling segregation. *CoRR* abs/2012.02086.

Chan, H.; Irfan, M. T.; and Than, C. V. 2020. Schelling models with localized social influence: a game-theoretic framework. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 240–248.

Chauhan, A.; Lenzner, P.; and Molitor, L. 2018. Schelling segregation with strategic agents. In *Proceedings of the 11th International Symposium on Algorithmic Game Theory (SAGT)*, 137–149.

Clark, W.; and Fossett, M. 2008. Understanding the social context of the Schelling segregation model. *Proceedings of the National Academy of Sciences* 105(11): 4109–4114.

Echzell, H.; Friedrich, T.; Lenzner, P.; Molitor, L.; Pappik, M.; Schöne, F.; Sommer, F.; and Stangl, D. 2019. Convergence and hardness of strategic Schelling segregation. In *Proceedings of the 15th International Conference on Web and Internet Economics (WINE)*, 156–170.

Elkind, E.; Fanelli, A.; and Flammini, M. 2020. Price of Pareto optimality in hedonic games. *Artificial Intelligence* 288: 103357.

Elkind, E.; Gan, J.; Igarashi, A.; Suksompong, W.; and Voudouris, A. A. 2019. Schelling games on graphs. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 266–272.

Immorlica, N.; Kleinberg, R.; Lucier, B.; and Zadimoghaddam, M. 2017. Exponential segregation in a two-dimensional Schelling model with tolerant individuals. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 984–993.

Kanellopoulos, P.; Kyropoulou, M.; and Voudouris, A. A. 2020. Modified Schelling games. In *Proceedings of the 13th International Symposium on Algorithmic Game Theory (SAGT)*, 241–256.

Koutsoupias, E.; and Papadimitriou, C. H. 1999. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, 404–413.

Pancs, R.; and Vriend, N. J. 2007. Schelling's spatial proximity model of segregation revisited. *Journal of Public Economics* 91(1): 1–24.

Pollicott, M.; and Weiss, H. 2001. The dynamics of Schelling-type segregation models and a nonlinear graph Laplacian variational problem. *Advances in Applied Mathematics* 27(1): 17–40.

Schelling, T. C. 1969. Models of segregation. *American Economic Review* 59(2): 488–493.

Schelling, T. C. 1971. Dynamic models of segregation. *Journal of Mathematical Sociology* 1(2): 143–186.

Young, H. P. 2001. *Individual Strategy and Social Structure: an Evolutionary Theory of Institutions*. Princeton University Press.

Zhang, J. 2004. A dynamic model of residential segregation. *Journal of Mathematical Sociology* 28(3): 147–170.