

Swap Stability in Schelling Games on Graphs*

Aishwarya Agarwal, Edith Elkind, Jiarui Gan, and Alexandros A. Voudouris

Department of Computer Science, University of Oxford

Abstract

We study a recently introduced class of strategic games that is motivated by and generalizes Schelling’s well-known residential segregation model. These games are played on undirected graphs, with the set of agents partitioned into multiple types; each agent either occupies a node of the graph and never moves away or aims to maximize the fraction of her neighbors who are of her own type. We consider a variant of this model that we call swap Schelling games, where the number of agents is equal to the number of nodes of the graph, and agents may *swap* positions with other agents to increase their utility. We study the existence, computational complexity and quality of equilibrium assignments in these games, both from a social welfare perspective and from a diversity perspective.

1 Introduction

Segregation is observed in many communities; people tend to group together on the basis of politics, religion, or socioeconomic status. This phenomenon has been extensively documented in residential metropolitan areas, where people may select where to live based on the racial composition of the neighborhoods. To formalize and study how the motives of individuals lead to residential segregation, Schelling [1969, 1971] proposed the following simple, yet elegant model. There are two types of agents who are to be placed on a line or a grid. An agent is happy with her location if at least a fraction $\tau \in (0, 1]$ of the agents within a certain radius are of the same type as her. Happy agents do not want to move, but unhappy agents are willing to do so in hopes of improving their current situation. Schelling described a dynamic process where at each step unhappy agents jump to random open positions or swap positions with other randomly selected agents, and showed that it can lead to a completely segregated placement, even if the agents themselves are tolerant of mixed neighborhoods ($\tau < 1/2$).

Over the years, Schelling’s work became very popular among researchers in Sociology and Economics, who proposed and studied numerous variants of his model, mainly via agent-based simulations; see the paper of Clark and Fossett [2008] and references therein for examples of this approach. Variants of the model have also been rigorously analyzed in a series of papers [Young, 2001; Zhang, 2004a; Brandt *et al.*, 2012; Barmpalias *et al.*, 2014; Bhakta *et al.*, 2014; Barmpalias *et al.*, 2015; Immorlica *et al.*, 2017], which showed that the random behavior of the agents leads with high probability to the formation of large monochromatic regions.

While all these papers focused on settings where the agents’ behavior is random, it is more realistic to assume instead that the agents are *strategic* and move only when they have an opportunity to improve their situation. So far, only a few papers have followed such a game-theoretic approach. In particular, Zhang [2004b] considered a game where the agents optimize a single-peaked utility function, and very recently, Chauhan *et al.* [2018], Elkind *et al.* [2019] and Echzell *et al.* [2019] studied strategic settings

*This work has been supported by the European Research Council (ERC) under grant number 639945 (ACCORD), and by the EPSRC International Doctoral Scholars Grant EP/N509711/1.

that are closer to the original model of Schelling, but capture more than two agent types and richer graph topologies.

In particular, Chauhan *et al.* [2018] study a setting with two types of agents, who have preferred locations, and can either swap with other agents or jump to empty positions. For a given tolerance threshold $\tau \in (0, 1]$, each agent’s primary goal is to maximize the fraction of her neighbors that are of her own type as long as this fraction is below τ (with all fractions above τ being equally good); her secondary goal is to be as close as possible to her preferred location. For both types of games (swap and jump), Chauhan *et al.* identify values of τ for which the best response dynamics of the agents leads to an equilibrium when the topology is a ring or a regular graph. Echzell *et al.* [2019] strengthen these results and extend them to more than two agent types, as well as study the complexity of computing assignments that maximize the number of happy agents.

Elkind *et al.* [2019] consider a similar model with k types; however, they treat agents’ location preferences differently from Chauhan *et al.* Namely, in their model each agent is either stubborn (i.e., has a preferred location and is unwilling to move) or strategic (i.e., aims to maximize the fraction of her neighbors that are of her own type; this corresponds to setting $\tau = 1$ in the model of Chauhan *et al.*). They focus on jump games, i.e., games where agents may jump to empty positions, and analyze the existence and complexity of computing Nash equilibria, as well as prove bounds on the price of anarchy Koutsoupias and Papadimitriou [1999] and the price of stability Anshelevich *et al.* [2008].

1.1 Our Contribution

We combine the two approaches by considering swap games in the model of Elkind *et al.* [2019]. That is, we assume that the number of agents is equal to the number of nodes in the topology, and two agents can swap locations if each of them prefers the other agent’s location to her own. We begin by studying the existence of equilibrium assignments. While such assignments exist for highly structured topologies, we show that they may fail to exist in general, even for simple topologies such as trees. Moreover, we show that deciding whether an equilibrium exists is NP-complete. We also study the quality of assignments in terms of their social welfare: we prove bounds on the price of anarchy and the price of stability for many interesting cases, and show that computing an assignment with high social welfare is NP-complete; the latter result complements the result of Elkind *et al.* in that it applies to the case where the number of agents equals the number of nodes in the topology.

Given that the goal of Schelling’s work was to study integration, it is natural to ask what level of integration can be achieved at equilibrium. There is a number of integration indices that have been proposed for this purpose (see, e.g., the survey of Massey and Denton [1988]). However, many of the indices defined in the literature are formulated for settings where the topology is highly regular and there are only two agent types, and it is not immediately clear how to adapt them to our model. We therefore focus on a simple index, which we call the *degree of integration*, that is inspired by the work of Lieberman and Carter [1982] and admits a natural interpretation in our context. This index counts the number of agents who are exposed to agents of other types, i.e., have at least one neighbor of a different type. We then study the price of anarchy and the price of stability with respect to this index: that is, we compare the value of our index in the best and worst equilibrium of our game to the optimal value of this index that can be achieved for a given instance. We note that, to the best of our knowledge, this is the first result of this type in the context of Schelling games: the previous work on integration in the Schelling model typically focused on evaluating a given integration index after some number of steps of the underlying dynamical process, and did not ask what level of integration can be achieved if the agents were non-strategic. We obtain strong negative results: it turns out that even the best equilibria are often much less diverse than the maximally diverse assignments. However, when the topology is a line, the price of stability with respect to our index can be bounded by a small constant. We also show that maximizing diversity is computationally hard.

1.2 Further Related Work

As mentioned above, Schelling’s model has been studied extensively both empirically and theoretically. For an introduction to the model and a survey of its many variants, we refer the reader to the book of Easley and Kleinberg [2010], and the papers by Brandt *et al.* [2012] and Immorlica *et al.* [2017]. Besides the closely related papers by Chauhan *et al.* [2018], Elkind *et al.* [2019] and Echzell *et al.* [2019], another work that is similar in spirit is a recent paper by Massand and Simon [2019], who study swap stability in games where a set of items is to be allocated among agents who are connected via a social network, so that each agent gets one item, and her utility depends on the items she and her neighbors in the network get; however, their results are not directly applicable to our setting. Also, Schelling games share a number of properties with hedonic games Drèze and Greenberg [1980]; Bogomolnaia and Jackson [2002], and in particular, with fractional hedonic games Aziz *et al.* [2019] and hedonic diversity games Bredereck *et al.* [2019]. However, a fundamental difference between hedonic games and Schelling games is that in the former agents form pairwise disjoint coalitions, while in the latter the neighborhoods of different nodes of the topology may overlap.

2 Preliminaries

A k -swap game is given by a set N of $n \geq 2$ agents partitioned into $k \geq 2$ pairwise disjoint types T_1, \dots, T_k , and an undirected simple connected graph $G = (V, E)$ with $|V| = n$, called the *topology*. We often identify types with colors: e.g., in a 2-swap game each agent is either red (T_1) or blue (T_2). The agents are also classified as either *strategic* or *stubborn*. We denote by R the set of strategic agents and by S the set of stubborn agents, so that $R \cup S = N$. Stubborn agents never move away from the nodes they occupy, while a strategic agent aims to maximize her personal utility, and is willing to swap positions with other agents to achieve this.

Given an agent $i \in T_\ell$, we refer to all other agents in T_ℓ as *friends of i* and denote the set of i ’s friends by $F_i = T_\ell \setminus \{i\}$. Each agent i occupies some node $v_i \in V$ of the topology G so that $v_i \neq v_j$ for every pair of agents $i \neq j$. The vector $\mathbf{v} = (v_1, \dots, v_n)$ that lists the locations of all agents is called an *assignment*. Given an assignment \mathbf{v} , we denote by $\pi_v(\mathbf{v})$ the agent that occupies node $v \in V$, that is, $\pi_{v_i}(\mathbf{v}) = i$.

Given an assignment \mathbf{v} , let $N_i(\mathbf{v}) = \{j \neq i : \{v_i, v_j\} \in E\}$ be the set of neighbors of agent i . The utility u_i of a stubborn agent $i \in S$ is independent of the assignment; e.g., we can set $u_i(\mathbf{v}) = 0$ for each $i \in S$. The utility of a strategic agent $i \in R$ for assignment \mathbf{v} is

$$u_i(\mathbf{v}) = \frac{|N_i(\mathbf{v}) \cap F_i|}{|N_i(\mathbf{v})|}.$$

Observe that, since $|V| = n$, every node is occupied by some agent, and therefore $N_i(\mathbf{v}) \neq \emptyset$ for every $i \in N$.

For every assignment \mathbf{v} , let $\mathbf{v}^{i \leftrightarrow j}$ be the assignment that is obtained from \mathbf{v} by swapping the positions of agents i and j : $v_\ell^{i \leftrightarrow j} = v_\ell$ for every $\ell \in N \setminus \{i, j\}$, $v_i^{i \leftrightarrow j} = v_j$ and $v_j^{i \leftrightarrow j} = v_i$. Agents i and j swap positions if and only if they both strictly increase their utility: $u_i(\mathbf{v}^{i \leftrightarrow j}) > u_i(\mathbf{v})$ and $u_j(\mathbf{v}^{i \leftrightarrow j}) > u_j(\mathbf{v})$. Clearly, agents of the same type cannot both increase their utilities by swapping, so swaps always involve agents of different types. An assignment \mathbf{v} is an *equilibrium* if no pair of agents i, j want to swap positions. That is, \mathbf{v} is an equilibrium if and only if for every $i, j \in R$ we have $u_i(\mathbf{v}) \geq u_i(\mathbf{v}^{i \leftrightarrow j})$ or $u_j(\mathbf{v}) \geq u_j(\mathbf{v}^{i \leftrightarrow j})$. We denote the set of all equilibrium assignments of the k -swap game \mathcal{G} by $\text{EQ}(\mathcal{G})$.

For every assignment \mathbf{v} , we define two benchmarks that aim to capture, respectively, the agents’ happiness and the societal diversity. The first one is the well-known *social welfare*, defined as the total utility of all strategic agents:

$$\text{SW}(\mathbf{v}) = \sum_{i \in R} u_i(\mathbf{v}).$$




Figure 1: The topology of the 2-swap game considered in the proof of Theorem 3.1, and an assignment that corresponds to the last case in the analysis; it is not an equilibrium since the red agent at node α and the blue agent at node γ_2 would like to swap.

Our second benchmark is the *degree of integration*: we say that an agent is *exposed* if she has at least one neighbor of a different type, and count the number of exposed agents:

$$DI(\mathbf{v}) = |\{i \in R : N_i(\mathbf{v}) \setminus F_i \neq \emptyset\}|.$$

Note that we ignore the stubborn agents in the definitions of our benchmarks, as their utility is always the same and they never want to move somewhere else.

For $f \in \{\text{SW}, \text{DI}\}$, let $\mathbf{v}_f^*(\mathcal{G})$ be the *optimal* assignment in terms of the benchmark f for a given game \mathcal{G} . The *price of anarchy* (PoA) in terms of the benchmark f is the worst-case ratio (over all k -swap games \mathcal{G} such that $\text{EQ}(\mathcal{G}) \neq \emptyset$) between the optimal performance (among all assignments) and the performance of the *worst* equilibrium assignment. Similarly, the *price of stability* (PoS) in terms of f is the worst-case ratio between the optimal performance and the performance of the *best* equilibrium:

$$\text{PoA}_f = \sup_{\mathcal{G}: \text{EQ}(\mathcal{G}) \neq \emptyset} \sup_{\mathbf{v} \in \text{EQ}(\mathcal{G})} \frac{f(\mathbf{v}_f^*(\mathcal{G}))}{f(\mathbf{v})},$$

$$\text{PoS}_f = \sup_{\mathcal{G}: \text{EQ}(\mathcal{G}) \neq \emptyset} \inf_{\mathbf{v} \in \text{EQ}(\mathcal{G})} \frac{f(\mathbf{v}_f^*(\mathcal{G}))}{f(\mathbf{v})}.$$

For readability, we refer to the quantity PoA_{SW} as the *social price of anarchy* and to PoA_{DI} as the *integration price of anarchy*, and use similar language for the price of stability.

3 Existence of Equilibria

We begin by discussing the existence of equilibria in swap games. The work of Echzell *et al.* [2019] implies that at least one equilibrium assignment exists when the topology is a regular graph. Furthermore, using a potential function similar to the one used by Elkind *et al.* [2019] for jump games, we can show that equilibria always exist when the topology is a graph of maximum degree 2; we omit the details. Our first result is a proof of non-existence of equilibria for every $k \geq 2$ for general topologies.

Theorem 3.1. *For every $k \geq 2$, there exists a k -swap game that does not admit an equilibrium assignment, even when all agents are strategic and the topology is a tree.*

Proof. We start with the easiest case of $k = 2$. Consider a 2-swap game with 10 strategic agents: 5 red agents and 5 blue agents. The topology is a tree with a root node α , which has three children nodes (set B), each of which has two children of its own (set Γ); see Figure 1. Suppose for the sake of contradiction that this game admits an equilibrium assignment \mathbf{v} .

Since $|B| = 3$ and there are only two types of agents, at least two nodes in B , say β_1 and β_2 , must be occupied by agents of the same type, say red. Now assume that nodes γ_1 (a child of β_1) and γ_2 (a child of

β_2) are occupied by blue agents. Then the red agent $\pi_{\beta_1}(\mathbf{v})$ and the blue agent $\pi_{\gamma_2}(\mathbf{v})$ can swap positions to increase their utility from strictly smaller than 1 and 0 to 1 and positive, respectively. Therefore, for at least one of these nodes (say, β_1) it must be the case that both of its children are occupied by red agents; as there are only five red agents, it follows that at least one of the children of β_2 , say γ_2 , is occupied by a blue agent.

If node α is occupied by a blue agent, then the red agent $\pi_{\beta_1}(\mathbf{v})$ and the blue agent $\pi_{\gamma_2}(\mathbf{v})$ can both increase their utility by swapping. Hence, node α must be occupied by a red agent (see Figure 1). However, this assignment is not an equilibrium either, since the red agent $\pi_{\alpha}(\mathbf{v})$ and the blue agent $\pi_{\gamma_2}(\mathbf{v})$ have an incentive to swap.

For $k \geq 3$, consider a k -swap game with $n = k(k^2 - 2)$ agents such that there are $k^2 - 2$ agents of type T_ℓ , for every $\ell \in [k]$. The topology is a tree whose nodes are distributed over three layers, just like in the case $k = 2$. Specifically, there is a root node α , which has a set B of $k(k - 1) - 1$ children. Each node in $\beta \in B$ has a set Γ_β of k children leaf nodes; let $\Gamma = \bigcup_{\beta \in B} \Gamma_\beta$. Next, we will argue about the structure of assignments that cannot be equilibria.

Lemma 3.2. *An assignment \mathbf{v} according to which any two nodes $\beta_1, \beta_2 \in B$ are occupied by agents of the same type T_x cannot be an equilibrium in case agents of some type T_y , $y \neq x$ simultaneously occupy nodes of Γ_{β_1} and Γ_{β_2} .*

Proof. Let \mathbf{v} be an assignment according to which nodes $\beta_1, \beta_2 \in B$ are occupied by agents of type T_x , and there exist nodes $\gamma_1 \in \Gamma_{\beta_1}$ and $\gamma_2 \in \Gamma_{\beta_2}$, which are occupied by agents of some type T_y , with $y \neq x$. Clearly, agent $\pi_{\beta_1}(\mathbf{v})$ has utility strictly less than 1, while agent $\pi_{\gamma_2}(\mathbf{v})$ has utility 0. Therefore, they would like to swap positions in order to increase their utility to 1 and positive, respectively. \square

Lemma 3.3. *An assignment \mathbf{v} according to which agents of every type occupy nodes of B cannot be an equilibrium.*

Proof. Let \mathbf{v} be an assignment according to which at least one agent of every type occupies some node of B . Without loss of generality, assume that the agent $\pi_{\alpha}(\mathbf{v})$ is of type T_x . We now distinguish between the following two cases.

- An agent of type T_x located at node $\beta \in B$ has a neighbor of type T_y , $y \neq x$ located at some node $\gamma \in \Gamma_\beta$. Then, by the assumption of the lemma, there exists at least one agent of type T_y located at some node $\beta' \in B \setminus \{\beta\}$, and therefore agents $\pi_{\alpha}(\mathbf{v})$ and $\pi_{\gamma}(\mathbf{v})$ would like to swap positions in order to increase their utility from strictly less than 1 and 0 to 1 and positive, respectively.
- For every agent of type T_x located at some node $\beta \in B$, all agents occupying the nodes of Γ_β are of type T_x . Since α is occupied by an agent of type T_x , there are $k^2 - 3 = (k - 1)(k + 1) - 2$ other agents of type T_x that can completely fill up at most $k - 2$ subtrees of α (since each of them consists of $k + 1$ nodes). Consequently, there are at least $k - 1$ agents of type T_x located at leaf nodes in other subtrees of α .

Now, assume that one of these agents of type T_x occupies a node $\gamma \in \Gamma_\beta$ such that $\beta \in B$ is occupied by an agent of type T_y , with $y \neq x$. We will now argue that there must exist another agent of type T_y located at some node $\beta' \in B \setminus \{\beta\}$. Assume otherwise that there is no such agent. Then, the remaining $k(k - 1) - 2$ agents of type T_y occupy leaf nodes of the tree. By Lemma 3.2, such agents located in different subtrees of α cannot be connected to agents of the same type. Hence, to cover all these $k(k - 1) - 2$ agents of type T_y , agents of $k - 1$ different types need to occupy the root nodes of the corresponding subtrees of α . However, there are only $k - 2$ types left (different than T_x and T_y). Consequently, there must exist another agent of type T_y that occupies some node $\beta' \in B \setminus \{\beta\}$. As a result, agents $\pi_{\gamma}(\mathbf{v})$ and $\pi_{\beta'}(\mathbf{v})$ have incentive to swap positions in order to increase their utility from 0 and strictly less than 1 to positive and 1, respectively.

This completes the proof of the lemma. \square

Lemma 3.4. *An assignment according to which there exists a type T_ℓ such that no agent of this type appears at nodes of B cannot be an equilibrium.*

Proof. Let \mathbf{v} be an assignment according to which no agent of type T_ℓ appears at the nodes of B . We first deal with the case $k \geq 4$. Observe that there are at least $k^2 - 3$ agents of type T_ℓ that must occupy nodes of Γ ; one agent of type T_ℓ may occupy α . By Lemma 3.2, agents of the same type that are located in different subtrees of α cannot be connected to agents of the same type. Hence, to cover all these $k^2 - 3$ agents of type T_ℓ , agents of k different types must occupy the root nodes of the corresponding subtrees of α , which is impossible.

For $k = 3$, if α is not occupied by an agent of type T_ℓ , then all $k^2 - 2 = 7$ agents of this type must occupy nodes of Γ , and the same argument as above leads to a contradiction. Hence, assume that $\pi_\alpha(\mathbf{v})$ is of type T_ℓ . Since $|B| = 5$ and no agent of type T_ℓ appears at the nodes of B , at least three nodes of B must be occupied by agents of the same type. Let $B = \{\beta_1, \dots, \beta_5\}$, and assume that nodes β_1, β_2 and β_3 are occupied by agents of type T_0 . If an agent of type T_ℓ occupies some node $\gamma \in \Gamma_\beta$ for any $\beta \in \{\beta_1, \beta_2, \beta_3\}$, then agent $\pi_{\beta'}(\mathbf{v})$, $\beta' \in \{\beta_1, \beta_2, \beta_3\} \setminus \{\beta\}$ has incentive to swap positions with agent $\pi_\gamma(\mathbf{v})$ to increase both of their utilities from strictly less than 1 and 0 to 1 and positive. Hence, no agent of type T_ℓ can be located at the nodes of $\Gamma_{\beta_1} \cup \Gamma_{\beta_2} \cup \Gamma_{\beta_3}$. Clearly, if agent $\pi_{\beta_4}(\mathbf{v})$ or $\pi_{\beta_5}(\mathbf{v})$ is of type T_ℓ , for the same reason, no agent of type T_ℓ can be located at the corresponding leaf nodes. Hence, both β_4 and β_5 must be occupied by agents of the third type T_1 and all leaf nodes $\Gamma_{\beta_4} \cup \Gamma_{\beta_5}$ must be occupied by the remaining 6 agents of type T_ℓ . However, this clearly cannot be an equilibrium assignment, since agent $\pi_{\beta_4}(\mathbf{v})$ would like to swap with any agent in Γ_{β_5} . \square

By Lemmas 3.3 and 3.4, we conclude that no assignment can be an equilibrium. \square

The topology used in the proof of Theorem 3.1 for the case $k = 2$ is utilized as a subgraph in the proof of the following theorem, to show that the problem of deciding whether an equilibrium exists is computationally hard.

Theorem 3.5. *For every $k \geq 2$, it is NP-complete to decide whether a given k -swap game admits an equilibrium.*

Proof. Membership in NP is immediate: we can verify whether a given assignment is an equilibrium by simply checking if there exists a pair of agents that would like to swap positions. To prove NP-hardness, we give a reduction from the CLIQUE problem, which is known to be NP-complete. An instance of CLIQUE consists of an undirected connected graph $H = (X, Y)$ and an integer λ ; it is a yes-instance if H contains a complete subgraph of size λ . Without loss of generality, we assume that $\lambda > 5$.

Given an instance $\langle H, \lambda \rangle$ of CLIQUE with $H = (X, Y)$, we will construct a 2-swap game as follows (the reduction can be extended to any $k > 2$ by adding stubborn agents of different types). Let d_v denote the degree of node v in H , and set $d_H = \max_{v \in X} d_v$.

- There are λ strategic red agents and $t = |X| + 5$ strategic blue agents; all other agents are stubborn, and will be defined in conjunction with the topology.
- The topology $G = (V, E)$ consists of three components G_1, G_2 and G_3 . These are connected to each other via stubborn agents, and are defined as follows:
 - To define $G_1 = (V_1, E_1)$, let W_v be a set of $2d_H - d_v + 2\lambda - 3$ nodes for each $v \in V$. Then, $V_1 = X \cup \bigcup_{v \in V} W_v$ and $E_1 = Y \cup \{\{v, w\} : v \in X, w \in W_v\}$. For every $v \in X$, d_H nodes of W_v are occupied by stubborn red agents, while the remaining $d_H - d_v + 2\lambda - 3$ nodes are occupied by stubborn blue agents. Observe that every node of G_1 has degree $d_1 = 2d_H + 2\lambda - 3$.

- The subgraph $G_2 = (A \cup B, E_2)$ is a complete bipartite graph with $|A| = \lambda - 5$ and $|B| = 4d_1$. Out of the $4d_1$ nodes of B , $2d_1 + 1$ nodes are occupied by stubborn red agents, while the remaining $2d_1 - 1$ nodes are occupied by stubborn blue agents.

Hence, a strategic red agent occupying a node of A has utility $\chi_r = \frac{2d_1+1}{4d_1} = \frac{1}{2} + \frac{1}{4d_1}$. Similarly, a strategic blue agent has utility $\chi_b = \frac{2d_1-1}{4d_1} = \frac{1}{2} - \frac{1}{4d_1}$.

- To define $G_3 = (V_3, E_3)$, let (V'_3, E'_3) be the graph used in the proof of Theorem 3.1, for which there is no equilibrium assignment; see Figure 1. For every node $v \in V'_3$ of degree 3, let Z_v be a set of $10d_1$ nodes such that $5d_1$ of these nodes are occupied by stubborn red agents, while the remaining $5d_1$ nodes are occupied by stubborn blue agents. For every $v \in V'_3$ of degree 1, let Z_v be a set of $10d_1$ nodes such that $5d_1$ of these nodes are occupied by stubborn red agents, while the remaining $5d_1$ nodes are occupied by stubborn blue agents. Then, $V_3 = V'_3 \cup_{v \in V'_3} Z_v$ and $E_3 = E'_3 \cup \{\{v, w\} : v \in V'_3, w \in Z_v\}$.

One can easily verify that the utility of a strategic agent (red or blue) occupying a node of G_3 is at least $\psi_0 = \frac{5d_1-1}{10d_1+1} > \frac{1}{2} - \frac{1}{4d_1}$ and at most $\psi_1 = \frac{5d_1+1}{10d_1+1} < \frac{1}{2} + \frac{1}{4d_1}$.

Now, assume that H has a clique of size λ , and let \mathbf{v} be the assignment in which the strategic red agents occupy the nodes of the clique, and the strategic blue agents occupy the remaining nodes. Each strategic red agent is connected to $\lambda - 1 + d_H$ other red agents (strategic and stubborn) in G_1 , and thus has utility

$$u = \frac{\lambda - 1 + d_H}{d_1} = \frac{d_H + \lambda - 1.5 + 0.5}{2d_H + 2\lambda - 3} \geq \frac{1}{2} + \frac{1}{2d_1}.$$

Clearly, since $u > \chi_r$ and $u > \psi_1$, no strategic red agent would be willing to swap positions with another strategic agent in G_2 or G_3 . By swapping positions with a blue agent within G_1 , a strategic red agent would still have at most $\lambda - 1 + d_H$ friends, and since every node in G_1 has the same degree, her utility cannot be improved. Hence, no strategic red agent has a profitable deviation, and \mathbf{v} is an equilibrium.

Conversely, assume that H does not contain a clique of size λ , and for the sake of contradiction also assume that there is an equilibrium assignment \mathbf{v} .

Suppose that some strategic red agents are located in G_1 . It cannot be the case that each of them is adjacent to $\lambda - 1$ other strategic red agents, as this would mean that the nodes they occupy form a clique of size λ . Hence, at least one of these agents, say agent i , is adjacent to at most $\lambda - 2$ strategic red agents. Since every node of G_1 has degree d_1 and every node is adjacent to d_H stubborn red agents, the utility of i is

$$u_i \leq \frac{d_H + \lambda - 2}{d_1} = \frac{d_H + \lambda - 1.5 - 0.5}{2d_H + 2\lambda - 3} = \frac{1}{2} - \frac{1}{2d_1}.$$

We have that $u_i < \chi_r$ and $u_i < \psi_0$, and hence agent i has incentive to move to G_2 or G_3 . On the other hand, the utility that a strategic blue agent j that is currently located in G_2 or G_3 can obtain by swapping positions with i is

$$u_j = 1 - u_i \geq \frac{1}{2} + \frac{1}{2d_1}.$$

Since $u_j > \chi_b$ and $u_j > \psi_1$, agent j also has an incentive to swap positions with agent i , and hence \mathbf{v} cannot be an equilibrium assignment. Therefore, no strategic red agent is located in G_1 .

Similarly, observe that $\chi_r > \psi_1$ and $\chi_b < \psi_0$, meaning that strategic red agents would prefer to be in G_2 , while strategic blue agents would prefer to be in G_3 . Thus, for \mathbf{v} to be an equilibrium assignment, it must be the case that all if a node of G_2 is not occupied by a stubborn agent, it is occupied by a strategic red agent. As a result, there are 5 strategic red and 5 strategic blue agents in G_3 . However, similarly to the proof of Theorem 3.1, we can argue that there is no equilibrium assignment for these agents in G_3 ; we omit the details here. Since we have exhausted all possibilities, it follows that if H does not have a clique of size λ , then there is no equilibrium assignment. \square




Figure 2: The topology and the equilibrium assignment of the lower bound instance in the proof of part (i) of Theorem 4.1. The big red square represents a clique whose nodes are occupied by red agents only.

4 Social Welfare

Here, we consider the efficiency of equilibrium assignments in terms of social welfare, and bound the social price of anarchy and stability for many interesting cases. We restrict our attention to games consisting of only strategic agents and such that there are at least two agents per type. Swap games with stubborn agents or strategic agents that are unique of their type can easily be seen to have unbounded social price of anarchy.¹

We start with the social price of anarchy of 2-swap games, and consider the general case (given the above restrictions) and the case where each type consists of the same number of agents.

Theorem 4.1. *The social price of anarchy of 2-swap games with only strategic agents is*

- (i) $\Theta(n)$ if there are at least two agents of each type, and
- (ii) between $921/448 \approx 2.0558$ and 4 if each type consists of the same number of agents.

Proof. We prove each part separately.

Part(i). For the lower bound, consider a 2-swap game with $n - 2$ red and 2 blue agents who are to be positioned on the nodes of a star topology. Then, the assignment according to which one of the blue agents occupies the central node is an equilibrium with social welfare $1 + \frac{1}{n-1} \leq 2$, while the optimal assignment is such that the central node is occupied by a red agent for a social welfare of $n - 3 + \frac{n-3}{n-1} \geq n - 3$.

For the upper bound, consider any 2-swap game with n agents such that there are $n_r \geq 2$ red and n_b blue agents, and let \mathbf{v} be any equilibrium assignment of this game. If there is any red agent ℓ with zero utility in \mathbf{v} , then it cannot be the case that this agent is connected to all blue agents. If this were the case, then since there is another red agent and the graph is connected, at least one blue agent must be connected to this red agent, get utility strictly less than 1, and have incentive to swap positions with agent ℓ so that they both increase their utility. Hence, in order for \mathbf{v} to be an equilibrium in the presence of ℓ getting zero utility, any blue agent not connected directly to ℓ must get utility 1 so that she does not want to swap with ℓ . Consequently, $\text{SW}(\mathbf{v}) \geq 1$. In the case where every agent has positive utility in \mathbf{v} , since the graph is connected, it must be the case that every agent get utility at least $1/n$, and therefore again $\text{SW}(\mathbf{v}) \geq 1$. Now the bound follows since the optimal social welfare is at most n .

Part (ii). For the lower bound, consider a 2-swap game with the following topology: there is a node α of degree $x + 1$ that is connected to x leaf nodes and to one node in a clique C of size $x - 1$. There is

¹For any $k \geq 2$, consider a k -swap game with a star topology and k types of agents such that there is only one red strategic agent, while the other types consist of at least two strategic agents and of some stubborn ones located at peripheral nodes. The assignment according to which the red agent occupies the center node is an equilibrium with 0 social welfare, while any assignment such that the center node is occupied by a non-red agent has positive social welfare.

an equilibrium \mathbf{v} where α is occupied by a red agent r , all leaf nodes are occupied by blue agents, and all nodes of C are occupied by red agents; see Figure 2. Hence,

$$\text{SW}(\mathbf{v}) = x - 1 + \frac{1}{x + 1} = \frac{x^2}{x + 1}.$$

On the other hand, for the assignment \mathbf{v}^* obtained from \mathbf{v} by swapping r with one of the blue agents we have

$$\text{SW}(\mathbf{v}^*) = 2x - 3 + \frac{x - 2}{x - 1} + \frac{x - 1}{x + 1}.$$

Hence, the price of anarchy is at least

$$\frac{2x^3 - x^2 - 5x + 2}{x^2(x - 1)},$$

an expression that takes its maximum value $667/324 \approx 2.05864$ at $x = 9$.

For the upper bound, consider a 2-swap game with $n = 2x$ agents such that there are $x \geq 2$ red and x blue agents. First, assume that some agents get zero utility in the equilibrium assignment \mathbf{v} . Observe that it cannot be the case that there exist agents of both types who have zero utility in \mathbf{v} . Indeed, if this was true for a non-adjacent red-blue pair, then these agents would have an incentive to swap and increase their utility from zero to 1. On the other hand, suppose this is true for an adjacent red-blue pair (r, b) . If both of r and b have other neighbors (besides r and b), then by swapping they can increase their utility from 0 to positive. Hence, suppose that r occupies a leaf node and is connected only to b . Then, since the graph is connected, there must exist a blue agent $b' \neq b$ with utility in $(0, 1)$ who would like to swap with r to increase both of their utilities from 0 to positive for r and from strictly less than 1 to 1 for b' .

Thus, assume that at least one blue agent has zero utility and all red agents have positive utility. We denote by B_0 the set of blue agents with zero utility, by R_1 the set of red agents with utility 1, and by $R_{<}$ the set of red agents whose utility is strictly less than 1. We have $|R_1| + |R_{<}| = x$, and each agent in B_0 is connected to all agents in $R_{<}$; otherwise a non-adjacent pair of agents $i \in R_{<}$, $j \in B_0$ would like to swap. If $|R_{<}| = 1$, then $|R_1| = x - 1$. The utility of the unique agent $i \in R_{<}$ is $u_i(\mathbf{v}) \geq \frac{1}{1+x}$, and thus

$$\text{SW}(\mathbf{v}) \geq x - 1 + \frac{1}{1 + x}.$$

Since the optimal social welfare is at most $n = 2x$ and $x \geq 2$, the price of anarchy is at most

$$\frac{2(1 + x)}{x} = 2 + \frac{2}{x} \leq 3.$$

Now, assume that $|R_{<}| \geq 2$. If $|B_0| \geq 2$, then any agent in B_0 would like to swap with any agent in $R_{<}$ to get positive utility (since each agent in $R_{<}$ is connected to all agents in B_0). Thus, in \mathbf{v} no agent $i \in R_{<}$ wants to swap with any agent in B_0 . Since each agent in B_0 is connected to all agents in $R_{<}$ and no other agent, the utility that agent i would get by agreeing to swap is $u_i(\mathbf{v}) = \frac{|R_{<}|-1}{|R_{<}|} \geq 1/2$, yielding

$$\text{SW}(\mathbf{v}) \geq |R_1| + |R_{>}|/2 \geq x/2.$$

Since the maximum welfare is $2x$, we can upper-bound the price of anarchy by 4. If $|B_0| = 1$, then if i is connected to another blue agent (not in B_0), her utility is at least $1/2$ for the same reason as before. Otherwise, i is connected to red agents only and the one agent in B_0 , so $u_i(\mathbf{v}) \geq 1/2$ (only one red agent in the worst case), which again yields an upper bound of 4 on the price of anarchy.

Next, we assume that all agents have positive utility. Since \mathbf{v} is an equilibrium, for every red-blue pair of agents it holds that at least one of them has no incentive to swap positions. Let (i, j) be a red-blue pair of agents, and assume that i does not want to swap. If i and j are *not* neighbors in \mathbf{v} , then it must be




Figure 3: The topology and the equilibrium assignment of the k -swap game considered in the proof of Theorem 4.2 for $k = 3$. Here, $T_1 = \text{red}$, $T_2 = \text{blue}$ and $T_3 = \text{green}$.

that $u_i(\mathbf{v}) \geq 1 - u_j(\mathbf{v})$ and hence $u_i(\mathbf{v}) + u_j(\mathbf{v}) \geq 1$. Otherwise, i and j are neighbors in \mathbf{v} . It may be the case that $u_i(\mathbf{v}) + u_j(\mathbf{v}) < 1$, in which case $u_i(\mathbf{v}) \in (0, 1)$ and $u_j(\mathbf{v}) \in (0, 1)$. Assume that the blue agent j has $x_r \geq 0$ red neighbors besides i , and $x_b \geq 1$ blue neighbors. Then, $u_j(\mathbf{v}) = \frac{x_b}{x_r + x_b + 1}$, and

$$\begin{aligned} u_i(\mathbf{v}) &\geq \frac{x_r}{x_r + x_b + 1} = 1 - u_j(\mathbf{v}) - \frac{1}{x_r + x_b + 1} \\ &\geq \frac{1}{2} - u_j(\mathbf{v}) \Leftrightarrow u_i(\mathbf{v}) + u_j(\mathbf{v}) \geq \frac{1}{2}, \end{aligned} \quad (1)$$

where the inequality follows since $\frac{1}{x_r + x_b + 1}$ is decreasing in $x_r \geq 0$ and $x_b \geq 1$. Therefore, in any case we have $u_i(\mathbf{v}) + u_j(\mathbf{v}) \geq \frac{1}{2}$, for every red-blue pair of agents i and j . Since there are x^2 distinct red-blue pairs, and each agent participates in exactly x such pairs, by summing over all these inequalities, we obtain $x \cdot \text{SW}(\mathbf{v}) \geq \frac{1}{2}x^2$ and therefore $\text{SW}(\mathbf{v}) \geq \frac{1}{2}x$. The bound then follows since the maximum possible social welfare is $n = 2x$. \square

We remark that even though the upper bound of 4 for the case where each type consists of the same number of agents (and every agent has positive utility at equilibrium) is probably not tight, one cannot expect to improve it using the same technique. In particular, to prove this upper bound, we focused on an arbitrary pair of agents (i, j) , and used the equilibrium definition, according to which at least one of these agents does not want to swap positions. Then, we were able to show that $u_i(\mathbf{v}) + u_j(\mathbf{v}) \geq 1/2$. We now argue that this inequality is actually tight, and hence to improve the upper bound one needs to argue in more detail about the structure of the equilibrium. Consider a variant of the topology depicted in Figure 2 in which every leaf node is connected to another leaf node. Then, for the depicted assignment \mathbf{v} , the red agent $\pi_{\alpha}(\mathbf{v})$ has utility $\frac{1}{x+1}$, while any blue agent has utility exactly $1/2$. Hence, the sum of the utility of the red agent $\pi_{\alpha}(\mathbf{v})$ and the utility of any blue agent is almost $1/2$ as x becomes large.

We continue by showing that, surprisingly, for three types or more, the social price of anarchy can be unbounded, even for the special case of equal number of agents per type.

Theorem 4.2. *For every $k \geq 3$, the social price of anarchy of k -swap games can be unbounded, even when there is an equal number of strategic agents per type.*

Proof. Consider a k -swap game with $n = 2k$ agents such that there are exactly two agents of each of the $k \geq 3$ types T_1, \dots, T_k . The topology G consists of k nodes $\{\alpha_1, \dots, \alpha_k\}$ that form a cycle, and each node α_ℓ , $\ell \in [k]$ is also connected to an auxiliary node β_ℓ ; see Figure 3 for the topology and the equilibrium assignment discussed in the following for $k = 3$.

Let \mathbf{v} be the assignment according to which node α_ℓ is occupied by an agent of type T_ℓ , while node β_ℓ is occupied by an agent of type $T_{\ell+1}$, where the subscripts are mod ℓ . This assignment is clearly an equilibrium since there exists no pair of agents that would like to swap positions, and every agent has zero utility. In particular, observe that agent $\pi_{\alpha_{\ell+1}}(\mathbf{v})$ of type $T_{\ell+1}$ would like to move only to node α_ℓ in order to connect with the agent $\pi_{\beta_\ell}(\mathbf{v})$ who is also of type $T_{\ell+1}$. However, the agent $\pi_{\alpha_\ell}(\mathbf{v})$ of type




Figure 4: The topology and the best equilibrium assignment for the 2-swap games considered in the proof of Theorem 4.3.

T_ℓ has no incentive to move to node $\alpha_{\ell+1}$ since the other agent of type T_ℓ is at node $\beta_{\ell-1}$. Now, consider agent $\pi_{\beta_\ell}(\mathbf{v})$ of type $T_{\ell+1}$ who is connected only to an agent of type T_ℓ located at α_ℓ . The only agent that would like to swap positions with $\pi_{\beta_\ell}(\mathbf{v})$ is $\pi_{\beta_{\ell-1}}(\mathbf{v})$ who is of type T_ℓ . However, such a swap cannot increase the utility of $\pi_{\beta_\ell}(\mathbf{v})$ since agent $\pi_{\alpha_{\ell-1}}(\mathbf{v})$ is of type $T_{\ell-1} \neq T_{\ell+1}$. Therefore, \mathbf{v} is an equilibrium with $\text{SW}(\mathbf{v}) = 0$.

On the other hand, consider the assignment \mathbf{v}^* according to which nodes α_ℓ and β_ℓ are occupied by the two agents of type T_ℓ , for every $\ell \in [k]$. Since every agent has now positive utility, $\text{SW}(\mathbf{v}^*) > 0$, and the social price of anarchy is unbounded. \square

Next, we turn our attention to the social price of stability and show a constant lower bound for 2-swap games.

Theorem 4.3. *The social price of stability of 2-swap games is at least $4/3$.*

Proof. Let $x \geq 3$ be a parameter, and consider a 2-swap game with $2x + 1$ red and $2x + 1$ blue agents. The topology is a tree with a root node α , which is connected to two nodes β and γ , as well as to a set A of $2x - 1$ leaf nodes. Moreover, node β is connected to a set B of x leaf nodes, and node γ is connected to a set Γ of x more leaf nodes. The topology and the best equilibrium assignment (which we discuss below) are depicted in Figure 4.

We will now argue about the structure of any equilibrium for this particular swap game. Without loss of generality, we assume that the root node α is occupied by a red agent, and switch between a few cases depending on the number of blue agents that occupy the leaf nodes of set A that are directly connected to α .

First, assume that there are at least $x + 1$ blue agents at the nodes of set A . Then, there can be at most $x - 2$ red agents at the nodes of A , which means that the remaining at least $x + 2$ red agents need to occupy nodes of the β - and γ -subtrees. Since any of these subtrees have a total of $\ell + 1$ nodes, at least one of these red agents, say agent i , must be connected to at least one blue agent. Clearly, such an assignment cannot be an equilibrium since agent i and any of the blue agents at the nodes of A have incentive to swap positions to increase their utilities from strictly smaller than 1 and 0 to 1 and positive, respectively.

Second, assume that there are exactly x blue agents at the nodes of set A , and hence the remaining $x - 1$ nodes of A are occupied by red agents. Then, it is easy to verify that the only equilibrium

assignment \mathbf{v}_1 (up to symmetries) is such that all nodes of the β -subtree are occupied by red agents, and all nodes of the γ -subtree are occupied by blue agents. The social welfare of this equilibrium is

$$\text{SW}(\mathbf{v}_1) = 3x + \frac{x}{x+1} + \frac{x}{2x+1} \leq 3x + 2.$$

Third, assume that the number of blue agents at the nodes of set A is between 1 and $x - 1$. Then, there are at least x red agents at the nodes of A . Since any of the β - and γ -subtrees have a total of $x + 1$ nodes, at least one of the remaining at most x red agents, say agent i , must necessarily be connected to some blue agent. As in the first case, such an assignment cannot be an equilibrium since agent i and any of the blue agents at the nodes of A have incentive to swap positions to increase their utilities from strictly smaller than 1 and 0 to 1 and positive, respectively.

Finally, assume that all nodes of A are occupied by red agents, and there is only one remaining red agent i , who will inevitably be connected to some blue agents. No assignment \mathbf{v}_2 according to which i occupies a leaf node of B (or Γ) can be an equilibrium, since i and the blue agent $\pi_\gamma(\mathbf{v}_2)$ (or $\pi_\beta(\mathbf{v}_2)$) have incentive to swap positions and increase their utilities from 0 and strictly smaller than 1 to positive and 1, respectively. Hence, in any equilibrium assignment \mathbf{v}_2 , agent i occupies either node β or node γ . The social welfare is

$$\text{SW}(\mathbf{v}_2) = 3x - 1 + \frac{2x}{2x+1} + \frac{1}{x+1} + \frac{x}{x+1} \leq 3x + 1.$$

Now consider the assignment \mathbf{v}^* according to which the red agents occupy node α , all nodes of A , and one node of B , while all other nodes are occupied by blue agents. The social welfare of this assignment is

$$\text{SW}(\mathbf{v}^*) = 4x - 1 + \frac{2x-1}{2x+1} + \frac{x-1}{x+1} \geq 4x - 1.$$

Therefore, the social price of anarchy is at least $\frac{4x-1}{3x+2}$, which tends to $4/3$ as x tends to infinity. \square

For k -swap games with topology that is a δ -regular graph (in which all nodes have degree equal to δ), we show an upper bound of 1 on the social price of stability by exploiting a potential function, similar to the one defined by Chauhan *et al.* [2018] and Echzell *et al.* [2019] to show the existence of equilibria in such games.

Theorem 4.4. *The social price of stability in k -swap games with topology that is a δ -regular graph is 1.*

Proof. Echzell *et al.* [2019] showed that for k -swap games with a δ -regular topology,

$$\Phi(\mathbf{v}) = \sum_{i \in R} |N_i(\mathbf{v}) \setminus F_i|$$

is a potential minimization function. Using similar arguments, we can show that the complement,

$$\bar{\Phi}(\mathbf{v}) = \sum_{i \in R} |N_i(\mathbf{v}) \cap F_i|,$$

is a potential maximization function. Consider any pair of agents (i, j) such that i is of type T_x and j is of type T_y , with $y \neq x$. Since i and j swap positions if and only if they can both increase their utility, and since $N_i(\mathbf{v}) = N_j(\mathbf{v}) = \delta$ for any assignment \mathbf{v} , we have that

$$|N_i(\mathbf{v}) \cap F_i| < |N_i(\mathbf{v}^{i \leftrightarrow j}) \cap F_i| \quad \text{and} \quad |N_j(\mathbf{v}) \cap F_j| < |N_j(\mathbf{v}^{i \leftrightarrow j}) \cap F_j|$$

Any agent $\ell \in (N_i(\mathbf{v}) \cap F_i) \cup (N_j(\mathbf{v}) \cap F_j)$ has one less friend in $\mathbf{v}^{i \leftrightarrow j}$ than in \mathbf{v} , and hence

$$|N_\ell(\mathbf{v}^{i \leftrightarrow j}) \cap F_\ell| = |N_\ell(\mathbf{v}) \cap F_\ell| - 1.$$

On the other hand, any agent $\ell \in (N_i(\mathbf{v}^{i \leftrightarrow j}) \cap F_i) \cup (N_j(\mathbf{v}^{i \leftrightarrow j}) \cap F_j)$ has one more friend in $\mathbf{v}^{i \leftrightarrow j}$ than in \mathbf{v} , and hence

$$|N_\ell(\mathbf{v}^{i \leftrightarrow j}) \cap F_\ell| = |N_\ell(\mathbf{v}) \cap F_\ell| + 1.$$

For any other agent, the friends they have as neighbors have not changed. Therefore, we can now easily see that

$$\bar{\Phi}(\mathbf{v}^{i \leftrightarrow j}) - \bar{\Phi}(\mathbf{v}) > 0,$$

as desired.

Now, observe that by the definition of the utility of each strategic agent and the fact that the topology is δ -regular, for any assignment \mathbf{v} , we have that

$$\text{SW}(\mathbf{v}) = \sum_{i \in R} u_i(\mathbf{v}) = \sum_{i \in R} \frac{|N_i(\mathbf{v}) \cap F_i|}{|N_i(\mathbf{v})|} = \frac{1}{\delta} \cdot \sum_{i \in R} |N_i(\mathbf{v}) \cap F_i| = \frac{1}{\delta} \cdot \bar{\Phi}(\mathbf{v}). \quad (2)$$

Let \mathbf{v}^* be an optimal assignment. If \mathbf{v}^* is an equilibrium, then the social price of stability is 1. Otherwise, we let the strategic agents play and swap positions until they reach an equilibrium \mathbf{v} . Since $\bar{\Phi}$ is a potential maximization function, we have that $\bar{\Phi}(\mathbf{v}) \geq \bar{\Phi}(\mathbf{v}^*)$, and by (2), we obtain

$$\text{SW}(\mathbf{v}) = \frac{1}{\delta} \cdot \bar{\Phi}(\mathbf{v}) \geq \frac{1}{\delta} \cdot \bar{\Phi}(\mathbf{v}^*) = \text{SW}(\mathbf{v}^*),$$

and the bound follows by rearranging terms. \square

Next, we focus on the problem of computing assignments of high social welfare. Observe that whether the agents are allowed to pairwise swap positions or jump to empty nodes has no effect in the complexity of this problem, and hence we already know that it is NP-complete by the work of Elkind *et al.* [2019]. However, one of their main assumptions is that the topology is a graph with strictly more nodes than agents (so that there are empty nodes where the agents can jump to). Consequently, their proof does not cover our case, where the topology consists of a number of nodes that is exactly equal to the number of agents. The proof of our next theorem is fundamentally different and subsumes that of Elkind *et al.* [2019] for $k \geq 3$; for $k = 2$, we were unable to prove the hardness of the problem.

Theorem 4.5. *For every $k \geq 3$, given a rational number ξ , it is NP-complete to decide whether there exists an allocation that has social welfare at least ξ .*

Proof. Membership in NP is immediate: given an assignment, we can sum up the utilities of the strategic agents and check whether the social welfare is at least ξ . To prove NP-hardness, we give a reduction from an NP-complete variant of the min-cut problem with additional cardinality constraints on the size of the subsets, to which we refer as the EQUAL-MIN-CUT problem Garey *et al.* [1974]. An instance of EQUAL-MIN-CUT consists of a graph $H = (X, Y)$, two distinguished nodes $s, t \in X$, and an integer W . It is a yes-instance if and only if there exist disjoint subsets of nodes X_1 and X_2 such that $X_1 \cup X_2 = X$, $|X_1| = |X_2|$, $s \in X_1$, $t \in X_2$ and $|\{\{v, z\} \in Y : v \in X_1, z \in X_2\}| \leq W$. Without loss of generality, we assume that $|X|$ is an even number, and by convention we denote an edge $\{v, z\}$ as vz to simplify our notation.

Given an instance $\langle H, s, t, W \rangle$ of EQUAL-MIN-CUT with $H = (X, Y)$, we construct an instance of our social welfare maximization problem as follows:

- There are $|X|/2 - 1$ strategic red and $|X|/2 - 1$ strategic blue agents.
- The topology $G = (V, E)$ consists of H with additional nodes and edges. Let s_0 and t_0 be two auxiliary nodes, and define $X_0 = \{s, s_0, t, t_0\}$ and $Y_0 = Y \cup \{s_0v : sv \in Y\} \cup \{t_0v : tv \in Y\}$. Let $d_v = |e \in Y_0 : v \in e|$ for every $v \in X \setminus X_0$, and $d_0 = \max_{v \in X \setminus X_0} d_v$. Also, let Z_v be a set of $d_0 - d_v$ nodes for every $v \in X \setminus X_0$. Then, G is such that $V = X \cup \{s_0, t_0\} \cup_{v \in X_0} Z_v$ and $E = Y_0 \cup \{vz : v \in X \setminus X_0, z \in Z_v\}$. Observe that in G , every node $v \in X \setminus X_0$ has degree exactly equal to d_0 .

- The nodes s and s_0 are occupied by stubborn red agents, t and t_0 are occupied by stubborn blue agents, and all nodes in $\bigcup_{v \in X \setminus X_0} Z_v$ are occupied by stubborn green agents.
- Finally, let $\xi = \frac{2}{d_0}(|Y| - W)$.

For any assignment \mathbf{v} and node $v \in V$, let $\chi_v(\mathbf{v})$ denote the type of the agent occupying node v . We will show that the social welfare of \mathbf{v} is decreasing in the number of edges of Y that are occupied by agents of different types. We have

$$\begin{aligned} \text{SW}(\mathbf{v}) &= \sum_{v \in X \setminus X_0} \frac{|\{vz \in Y_0 : \chi_v(\mathbf{v}) = \chi_z(\mathbf{v})\}|}{d_0} \\ &= \frac{1}{d_0} \sum_{v \in X \setminus X_0} |\{vz \in Y_0 : \chi_v(\mathbf{v}) = \chi_z(\mathbf{v}), z \notin X_0\}| \\ &\quad + \frac{1}{d_0} \sum_{v \in X \setminus X_0} |\{vz \in Y_0 : \chi_v(\mathbf{v}) = \chi_z(\mathbf{v}), z \in X_0\}|. \end{aligned}$$

Since $\chi_s(\mathbf{v}) = \chi_{s_0}(\mathbf{v})$, $\chi_t(\mathbf{v}) = \chi_{t_0}(\mathbf{v})$, $vs \in Y \Leftrightarrow vs_0 \in Y_0$, and $vt \in Y \Leftrightarrow vt_0 \in Y'$, it follows that for every $v \in X \setminus X_0$,

$$\begin{aligned} |\{vz \in Y_0 : \chi_v(\mathbf{v}) = \chi_z(\mathbf{v}), z \in X_0\}| &= 2|\{vz \in Y_0 : \chi_v(\mathbf{v}) = \chi_z(\mathbf{v}), z \in \{s, t\}\}| \\ &= 2|\{vw \in Y : \chi_v(\mathbf{v}) = \chi_w(\mathbf{v}), w \in \{s, t\}\}|. \end{aligned}$$

Furthermore,

$$|\{vz \in Y_0 : \chi_v(\mathbf{v}) = \chi_z(\mathbf{v}), z \notin X_0\}| = |\{vw \in Y_0 : \chi_v(\mathbf{v}) = \chi_w(\mathbf{v}), w \notin \{s, t\}\}|$$

Therefore, we obtain

$$\begin{aligned} \text{SW}(\mathbf{v}) &= \frac{2}{d_0} |\{vz \in Y : \chi_v(\mathbf{v}) = \chi_z(\mathbf{v}), z \in X \setminus \{s, t\}\}| \\ &\quad + \frac{2}{d_0} |\{vz \in Y : \chi_v(\mathbf{v}) = \chi_z(\mathbf{v}), z \in \{s, t\}\}| \\ &= \frac{2}{d_0} |\{vz \in Y : \chi_v(\mathbf{v}) = \chi_z(\mathbf{v})\}| \\ &= \frac{2}{d_0} (|Y| - |\{vz \in Y : \chi_v(\mathbf{v}) \neq \chi_z(\mathbf{v})\}|). \end{aligned} \tag{3}$$

Now, assume that the input instance of EQUAL-MIN-CUT is a yes-instance, and let $X = X_1 \cup X_2$ be the satisfying partition. Let \mathbf{v} be such that the strategic red agents occupy the nodes of X_1 and the strategic blue agents occupy the nodes of X_2 . Then, by definition we have that $|\{vz \in Y : \chi_v(\mathbf{v}) \neq \chi_z(\mathbf{v})\}| = |\{vz \in Y : v \in X_1, z \in X_2\}| \leq W$, and by (3) we obtain

$$\text{SW}(\mathbf{v}) \geq \frac{2}{d_0}(|Y| - W) = \xi.$$

Conversely, assume that there exists an assignment \mathbf{v} with $\text{SW}(\mathbf{v}) \geq \xi = \frac{2}{d_0}(|Y| - W)$. Let X_1 consist of the nodes occupied by strategic red agents, and let X_2 consist of the nodes occupied by strategic blue agents. Then, $X_1 \cap X_2 = \emptyset$, and since there is an equal number of strategic red and blue agents, we also have that $|X_1| = |X_2|$. By (3), it is

$$|\{vz \in Y : \chi_v(\mathbf{v}) \neq \chi_z(\mathbf{v})\}| \leq W,$$

and consequently

$$|\{vz \in Y : \chi_v(\mathbf{v}) \neq \chi_z(\mathbf{v})\}| = |\{vz \in Y : v \in X_1, z \in X_2\}| \leq W,$$

as desired. □




Figure 5: The topology and the only possible equilibrium assignment for the 2-swap game considered in the proofs of Theorems 5.1 and 5.2. For k -swap games (Theorem 5.1) there are k identical subtrees, and in the worst equilibrium each subtree is filled by agents of a different type.

5 Degree of Integration

We now investigate whether equilibrium assignments can be diverse, by bounding the price of anarchy and stability in terms of the degree of integration; recall that this benchmark counts the number of agents who are exposed, i.e., have at least one neighbor of a different type. As in the previous section, we again focus on games with strategic agents only.

We start by showing that the integration price of anarchy of k -swap games is n/k , i.e., it scales linearly with the number of agents. This indicates that in the worst case agents of different types are highly segregated, but, as the number of types increases, equilibria become more diverse and the price of anarchy decreases.

Theorem 5.1. *For any $k \geq 2$, the integration price of anarchy of k -swap games with strategic agents only is at most n/k , and this bound is tight.*

Proof. For the upper bound, consider a k -swap game with n agents. By definition, the optimal degree of integration is at most n . Since the topology is a connected graph, in any assignment \mathbf{v} at least one agent per type must be exposed. Hence, $\text{DI}(\mathbf{v}) \geq k$, and the integration price of anarchy is at most n/k .

For the lower bound, consider a k -swap game with $n = kx + 1$ agents such that there are $x + 1$ agents of type T_1 and x agents of type T_ℓ for every $\ell \in [k] \setminus \{1\}$. The topology is a tree with root node α that has k children nodes β_1, \dots, β_k , each of which has $x - 1$ children leaf nodes of its own; see Figure 5 for an example of this topology for $k = 2$.

One can assign the agents to the nodes of the topology so that each agent is exposed; thus the maximum possible degree of integration is n . However, there is an equilibrium assignment \mathbf{v} in which α is occupied by an agent of type T_1 and for each $\ell \in [k]$ all nodes of the β_ℓ -subtree are occupied by agents of type T_ℓ . In \mathbf{v} only the agent in α and the agents in nodes β_ℓ , $\ell \in [k] \setminus \{1\}$, are exposed, yielding degree of integration $\text{DI}(\mathbf{v}) = k$, and the bound follows. \square

Next, we consider the integration price of stability. Using the same instance as in the proof of Theorem 5.1, we show a lower bound that depends linearly on the number of agents for the fundamental case of two agent types. This bound is tight by the previous theorem, and indicates that we cannot avoid ending up with equilibrium assignments in which the types are highly segregated, even in the best case.

Theorem 5.2. *The integration price of stability of 2-swap games with strategic agents only is at least $n/2$.*

Proof. Consider a 2-swap game with $x + 1$ red agents and x blue agents, for a total of $n = 2x + 1$ agents. The topology is the same as in Theorem 8: a tree consisting of a root node α with two children

nodes β_1 and β_2 , each of which has $x - 1$ children leaf nodes of its own (sets B_1 and B_2); see Figure 5. The optimal degree of integration is n . We will now argue that the unique equilibrium assignment \mathbf{v} (up to symmetries) is such that α and all nodes of the β_1 -subtree are occupied by red agents, while the nodes of the β_2 -subtree are occupied by blue agents. The degree of integration of \mathbf{v} is exactly 2, so the theorem follows.

Assume that agent $\pi_\alpha(\mathbf{v})$ is blue rather than red. We distinguish the following cases with regards to the agents occupying nodes β_1 and β_2 .

- Both $\pi_{\beta_1}(\mathbf{v})$ and $\pi_{\beta_2}(\mathbf{v})$ are of the same type. Assume that both of these agents are blue; as there are $x + 1$ red agents, there must be red agents at the leaf nodes of both B_1 and B_2 . But then agent $\pi_{\beta_2}(\mathbf{v})$ and some red agent occupying a node of B_1 can swap to increase their utility from strictly less than 1 and zero to 1 and positive, respectively. Hence, it must be the case that $\pi_{\beta_1}(\mathbf{v})$ and $\pi_{\beta_2}(\mathbf{v})$ are both red. Again, if there are blue agents at the leaf nodes of both B_1 and B_2 , the assignment is not an equilibrium, so the nodes of one of these subtrees (say, B_1) are all occupied by red agents, and the nodes of the other subtree are all occupied by blue agents. However, such an assignment cannot be an equilibrium since the blue agent $\pi_\alpha(\mathbf{v})$ and the red agent $\pi_{\beta_2}(\mathbf{v})$ both get zero utility and have an incentive to swap.
- $\pi_{\beta_1}(\mathbf{v})$ is red and $\pi_{\beta_2}(\mathbf{v})$ is blue. Since there are x red agents remaining, at least one of them must be in B_2 . But then she can swap positions with the blue agent $\pi_\alpha(\mathbf{v})$ so that they increase their utility from zero and $1/2$ to $1/2$ and 1, respectively.

Therefore, agent $\pi_\alpha(\mathbf{v})$ must be red. Similarly to the previous case, we observe that if $\pi_{\beta_1}(\mathbf{v})$ and $\pi_{\beta_2}(\mathbf{v})$ are both blue, there must be red agents in both B_1 and B_2 , and if $\pi_{\beta_1}(\mathbf{v})$ and $\pi_{\beta_2}(\mathbf{v})$ are both red, there must be blue agents in both B_1 and B_2 , which means that $\pi_{\beta_1}(\mathbf{v})$ and some agent in a node of B_2 would have an incentive to swap. Thus, one of $\pi_{\beta_1}(\mathbf{v})$ and $\pi_{\beta_2}(\mathbf{v})$ (say, $\pi_{\beta_1}(\mathbf{v})$) must be red and the other one must be blue. Then, if there is a blue agent in B_1 , by a counting argument there is also a red agent in B_2 , and these two agents would have an incentive to swap. Hence, \mathbf{v} is the only equilibrium assignment. \square

To develop better intuition for the integration price of anarchy and stability, we also consider the special case where the topology is a line. In this case, while the integration price of anarchy remains linear in n/k , the integration price of stability can be bounded by a small constant.

Theorem 5.3. *Consider a k -swap game with strategic agents only, at least two agents per type, and a line topology. The integration price of anarchy is at most $\frac{n}{2k-2}$, while the integration price of stability is at most $\frac{9}{4}$. Moreover, if the number of agents of each type grows with n , the integration price of stability is at most $\frac{3}{2} + o(1)$. All these bounds are tight.*

Proof. Let the topology be a line, with nodes $1, \dots, n$ connected in this order.

Price of anarchy. For the upper bound, consider an equilibrium assignment \mathbf{v} . For each type T_i let ℓ_i be the leftmost agent of type T_i and let r_i be the rightmost agent of this type. If $v_{\ell_i} \neq 1$, then ℓ_i has a neighbor to the left who belongs to a different type; similarly, if $v_{r_i} \neq n$ then r_i has a neighbor to the right who belongs to a different type. Since nodes 1 and n are occupied by exactly two agents, it follows that $\text{DI}(\mathbf{v}) \geq 2k - 2$. As at most n agents are exposed in any assignment, the bound follows.

To see that this bound is tight, it suffices to consider a k -swap game with s agents per type, for some $s \geq 2$. We can create s identical blocks of agents, with each block containing exactly one agent of each type, and place them on the line one after the other, so that every agent is exposed. However, there is also an equilibrium where agents are placed in monochromatic blocks of size s , so that only $2k - 2$ agents are exposed.

Price of stability. We partition the agents of each type into blocks of size 2 and 3, with at most one block of size 3 per type (that is, we create a block of size 3 if and only if the number of agents of that type is odd); let B_1, \dots, B_d be the resulting collection of blocks. Observe that any assignment where agents in each block are placed contiguously on the line is an equilibrium. Indeed, under any such assignment each agent has at least one neighbor of the same type, and no agent can move to a position where she would have two neighbors of her type; she cannot move to nodes 1 or n either, since agents at these nodes are unwilling to swap (they have utility 1).

It remains to explain how to place these blocks on the line to maximize the degree of integration. We do so greedily, from left to right. That is, we first pick some $i \in [k]$ such that $|T_i| \geq |T_j|$ for all $j \in [k]$, and place some block $B \subseteq T_i$ first; if $|B| = 2$, we assign the agents in B to nodes 1 and 2, and if $|B| = 3$, we assign the agents in B to nodes 1, 2, and 3. Now, suppose that ℓ blocks have been placed, so that the last occupied node is node r , and the agent there is of type T_x . For each $j \in [k]$, let t_j be the number of agents of type T_j who have not yet been placed. If $t_j = 0$ for all $j \in [k] \setminus \{x\}$, we complete the assignment by simply placing all the remaining agents of type T_x on the line. Otherwise, we pick an $i \in [k] \setminus \{x\}$ such that $t_i \geq t_j$ for all $j \in [k] \setminus \{x\}$, and place some block $B \subseteq T_i$ in positions $r + 1, \dots, r + |B|$.

Let us say that a type T_i is *dominant* if $|T_i| > n/2$. An easy inductive argument shows that if no type is dominant, then under this assignment we never place two blocks of the same type next to each other; the key observation is that if no type is dominant after ℓ blocks have been placed, this remains to be the case after $\ell + 2$ blocks have been placed, and hence if we still have at least two blocks to place, we have at least two types to choose from. In this case, the only agents who are not exposed are agents at nodes 1 and n as well as agents located at the middle of a block of size 3, i.e., at most $k + 2$ agents. Thus, the integration price of stability is at most $\frac{n}{n-k-2}$ in this case. Now, suppose that some type (say, type T_1) is dominant. If there are λ blocks of types T_2, \dots, T_k , then under our procedure we will first alternate between blocks of type T_1 and blocks of other types, and then place the remaining blocks of type T_1 (if any). Then at least 4λ agents will be exposed. On the other hand, at most $k - 1$ of these λ blocks are of size 3, so we have at most $2\lambda + k - 1$ agents of types T_2, \dots, T_k , and hence in any assignment at most $2(2\lambda + k - 1)$ agents of type T_1 can be exposed. Thus, the integration price of stability in this case is at most

$$\frac{3(2\lambda + k - 1)}{4\lambda} = \frac{3}{2} + \frac{3k - 3}{4\lambda}.$$

Since $\lambda \geq k - 1$, this quantity is at most $\frac{9}{4}$. Further, if we assume that the number of agents of each type grows with n , we have $\frac{3k-3}{4\lambda} = o(1)$, and the bound becomes $\frac{3}{2} + o(1)$.

To see that the bound $\frac{9}{4}$ on the integration price of stability is tight, consider a game with six red agents and three blue agents. In equilibrium, the three blue agents need to form a single block: if there is an isolated blue agent b , there is also a red agent r who is not a neighbor of b , but has another blue neighbor; b and r can then benefit from swapping. Thus, in equilibrium at most two blue agents—and hence at most two red agents—are exposed. However, we can also create three blocks of agents, with each block consisting of a single blue agent surrounded by two red agents, and place these three blocks consecutively on the line, so that each agent is exposed.

To see that the bound $\frac{3}{2}$ is tight if the number of agents of each type grows with n , consider an instance with $4s$ red agents and $2s$ blue agents, for some $s \in \mathbb{N}$. Arguing as above, we can see that in equilibrium the blue agents have to appear in blocks of size at least 2, so that each blue agent has at most one red neighbor. Hence, at most $2s$ red agents have a blue neighbor, and thus the number of exposed agents is at most $4s$. On the other hand, by placing agents in red-blue-red blocks, as described in the previous paragraph, we can ensure that all $6s$ agents are exposed. \square

Hence, for games with simple line topologies, integration can be achieved in equilibrium. However, when left to their own devices, the agents may end up in a very segregated configuration.

We conclude this section by studying the complexity of computing assignments (not necessarily equilibria) with high degree of integration. Unfortunately, it turns out that even this task is computationally intractable.

Theorem 5.4. *Given a k -swap game, computing an assignment in which every agent is exposed is NP-complete, even if $k = 2$ and all agents are strategic.*

Proof. The problem is clearly in NP; for a given assignment we can verify whether each of the n agents has at least one neighbor of a different type in time $O(n^2)$. For the NP-hardness proof, we give a reduction from the VERTEX COVER problem, which is known to be NP-complete. An instance of VERTEX COVER consists of an undirected graph $H = (X, Y)$ and an integer $\lambda \leq |X|$. It is a yes-instance if there exists a set $X' \subseteq X$ such that $|X'| \leq \lambda$ and $\{v, w\} \cap X' \neq \emptyset$ for every edge $\{v, w\} \in Y$. Such a set X' is called a vertex cover of H . Without loss of generality, we assume that H has no isolated vertices.

Given an instance $\langle H, \lambda \rangle$ of the VERTEX COVER problem with $H = (X, Y)$, we construct a 2-swap game as follows:

- We have $|X| + |Y| - \lambda$ red agents and λ blue agents, for a total of $n = |X| + |Y|$ agents.
- To construct the topology $G = (V, E)$, we start with the graph H . Then, for each edge $e = \{v, w\} \in Y$, we add a node s_e , and two edges connecting s_e to v and w . Let $Q = \{s_e : e \in Y\}$. Then, $V = X \cup Q$, $X \cap Q = \emptyset$, and $|V| = |X| + |Q| = |X| + |Y| = n$.

We show that H has a vertex cover of size at most λ if and only if there exists an assignment in which every agent is exposed.

First, suppose that there exists a vertex cover $X' \subseteq X$ of H of size at most λ ; by adding nodes, we can assume that $|X'| = \lambda$. Consider the assignment \mathbf{v} in which the nodes of X' are occupied by blue agents, while all other nodes of $V \setminus X'$ are occupied by red agents. In this assignment, every agent is exposed:

- For every blue agent i occupying a node $v_i \in X'$, since H has no isolated nodes, there must exist an edge $e \in Y$ such that $v_i \in e$, and hence v_i is connected to s_e , which is occupied by a red agent.
- For every red agent i occupying a node $v_i \in X \setminus X'$, since X' is a vertex cover, v_i must be connected to a node $z \in X'$, which is occupied by a blue agent.
- For every red agent i occupying a node $v_i = s_e \in Q$, since X' is a vertex cover, at least one of e 's endpoints must be in X' , which is occupied by a blue agent, and s_e is connected to both endpoints of e .

Conversely, suppose that \mathbf{v} is an assignment of the agents to the nodes of the topology such that every agent is exposed.

For each edge $e = \{v, w\} \in E$, let $\ell(e)$ be an arbitrary element of $\{v, w\}$. Let $X' = \{v \in X : \pi_v(\mathbf{v}) \text{ is blue}\}$, $X_Q = \{z \in X : z = \ell(e) \text{ for some } e \text{ such that } \pi_{s_e}(\mathbf{v}) \text{ is blue}\}$. Since there are $\lambda - |X'|$ nodes in Q that are occupied by blue agents, we have $|X_Q| \leq \lambda - |X'|$ and hence $|X' \cup X_Q| \leq \lambda$. We claim that $X' \cup X_Q$ is a vertex cover for H . Indeed, consider an arbitrary edge $e = \{v, w\} \in E$; we will argue that $e \cap (X' \cup X_Q) \neq \emptyset$. If $v \in X'$ or $w \in X'$, we are done. Otherwise, both v and w are occupied by red agents; since $\pi_{s_e}(\mathbf{v})$ is adjacent to an agent of a different type, it follows that s_e is occupied by a blue agent and $\ell(e) \in X_Q$. Hence one of v and w is in X_Q . This completes the proof. \square

6 Conclusions and Open Problems

We have studied Schelling games on graphs in which pairs of agents can deviate by swapping their locations. We considered questions related to the existence and the efficiency of equilibrium assignments, both from a social welfare perspective and from a diversity perspective.

While equilibria are known to exist in instances where the topology is highly structured, we showed that their existence is not guaranteed in general, and deciding whether a given swap game admits an equilibrium assignment is NP-complete. Even though we have implicitly assumed that the tolerance threshold of every agent is 1, and thus she is never truly happy unless she is connected to friends only, our proofs extend to other threshold values as well. For instance, one can verify that Theorem 3.1 for $k = 2$ holds for any $\tau \in (2/3, 1)$. A challenging open question is to completely characterize the topologies and the threshold values for which equilibria are guaranteed to exist, and also design efficient algorithms to compute equilibria when they exist.

We have introduced a new index for measuring the diversity of a given assignment, which we called the degree of integration. It would be interesting to explore the tradeoffs between diversity and social welfare: can we compute (equilibrium) assignments with a given degree of integration that maximize the social welfare? While our results indicate that this problem is hard for general topologies, one could hope to obtain approximate or parameterized algorithms, or focus on simple topologies. One can also investigate more ambitious diversity indices, e.g., by considering, for each agent, the number of other types she is exposed to.

References

- Elliot Anshelevich, Anirban Dasgupta, Jon M. Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.
- Haris Aziz, Florian Brandl, Felix Brandt, Paul Harrenstein, Martin Olsen, and Dominik Peters. Fractional hedonic games. *ACM Transactions on Economics and Computation*, 7(2):6:1–6:29, 2019.
- George Barmpalias, Richard Elwes, and Andrew Lewis-Pye. Digital morphogenesis via Schelling segregation. In *Proceedings of the 55th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 156–165, 2014.
- George Barmpalias, Richard Elwes, and Andrew Lewis-Pye. From randomness to order: unperturbed Schelling segregation in two or three dimensions. *CoRR*, abs/1504.03809, 2015.
- Prateek Bhakta, Sarah Miracle, and Dana Randall. Clustering and mixing times for segregation models on \mathbb{Z}^2 . In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 327–340, 2014.
- Anna Bogomolnaia and Matthew O. Jackson. The stability of hedonic coalition structures. *Games and Economic Behavior*, 38(2):201–230, 2002.
- Christina Brandt, Nicole Immorlica, Gautam Kamath, and Robert Kleinberg. An analysis of one-dimensional Schelling segregation. In *Proceedings of the 44th Symposium on Theory of Computing Conference (STOC)*, pages 789–804, 2012.
- R. Bredereck, E. Elkind, and A. Igarashi. Hedonic diversity games. In *Proceedings of the 18th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 565–573, 2019.

- Ankit Chauhan, Pascal Lenzner, and Louise Molitor. Schelling segregation with strategic agents. In *Proceedings of the 11th International Symposium on Algorithmic Game Theory (SAGT)*, pages 137–149, 2018.
- William Clark and Mark Fossett. Understanding the social context of the Schelling segregation model. *Proceedings of the National Academy of Sciences*, 105(11):4109–4114, 2008.
- Jacques H. Drèze and Joseph Greenberg. Hedonic coalitions: optimality and stability. *Econometrica*, 48(4):987–1003, 1980.
- David A. Easley and Jon M. Kleinberg. *Networks, Crowds, and Markets – Reasoning about a Highly Connected World*. Cambridge University Press, 2010.
- Hagen Echezell, Tobias Friedrich, Pascal Lenzner, Louise Molitor, Marcus Pappik, Friedrich Schöne, Fabian Sommer, and David Stangl. Convergence and hardness of strategic Schelling segregation. *CoRR*, abs/1907.07513, 2019.
- Edith Elkind, Jiarui Gan, Ayumi Igarashi, Warut Suksompong, and Alexandros A. Voudouris. Schelling games on graphs. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 266–272, 2019.
- Michael R. Garey, David S. Johnson, and Larry Stockmeyer. Some simplified np-complete problems. In *Proceedings of the 6th annual ACM Symposium on Theory of Computing (STOC)*, pages 47–63, 1974.
- Nicole Immorlica, Robert Kleinberg, Brendan Lucier, and Morteza Zadimoghaddam. Exponential segregation in a two-dimensional Schelling model with tolerant individuals. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 984–993, 2017.
- Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 404–413, 1999.
- Stanley Lieberson and Donna K. Carter. A model for inferring the voluntary and involuntary causes of residential segregation. *Demography*, 19(4):511–526, 1982.
- Sagar Massand and Sunil Simon. Graphical one-sided markets. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 492–498, 2019.
- Douglas S. Massey and Nancy A. Denton. The dimensions of residential segregation. *Social Forces*, 67(2):281–315, 1988.
- Thomas C. Schelling. Models of segregation. *American Economic Review*, 59(2):488–493, 1969.
- Thomas C. Schelling. Dynamic models of segregation. *Journal of Mathematical Sociology*, 1(2):143–186, 1971.
- H. Peyton Young. *Individual Strategy and Social Structure: an Evolutionary Theory of Institutions*. Princeton University Press, 2001.
- Junfu Zhang. A dynamic model of residential segregation. *Journal of Mathematical Sociology*, 28(3):147–170, 2004.
- Junfu Zhang. Residential segregation in an all-integrationist world. *Journal of Economic Behavior and Organization*, 54(4):533–550, 2004.