Tolerance is Necessary for Stability: Single-Peaked Swap Schelling Games

Davide Bilò[1] , **Vittorio Bilò**[2] , **Pascal Lenzner**[3] and **Louise Molitor**[3]
[1]University of Sassari, Italy
[2]University of Salento, Italy
[3]Hasso Plattner Institute, University of Potsdam, Germany
davide.bilo@uniss.it, vittorio.bilo@unisalento.it, pascal.lenzner@hpi.de, louise.molitor@hpi.de

Abstract

Residential segregation in metropolitan areas is a phenomenon that can be observed all over the world. Recently, this was investigated via game-theoretic models. There, selfish agents of two types are equipped with a monotone utility function that ensures higher utility if an agent has more same-type neighbors. The agents strategically choose their location on a given graph that serves as residential area to maximize their utility. However, sociological polls suggest that real-world agents are actually favoring mixed-type neighborhoods, and hence should be modeled via non-monotone utility functions. To address this, we study Swap Schelling Games with single-peaked utility functions. Our main finding is that tolerance, i.e., agents favoring fifty-fifty neighborhoods or being in the minority, is necessary for equilibrium existence on almost regular or bipartite graphs. Regarding the quality of equilibria, we derive (almost) tight bounds on the Price of Anarchy and the Price of Stability. In particular, we show that the latter is constant on bipartite and almost regular graphs.

1 Introduction

Residential segregation is defined as the physical separation of two or more groups into different neighborhoods [Massey and Denton, 1988]. It is pervasive in metropolitan areas, where large homogeneous regions inhabited by residents belonging to the same ethnic group emerged over time[1].

For more than five decades, residential segregation has been intensively studied by sociologists, as a high degree of segregation has severe consequences for the inhabitants of homogeneous neighborhoods. It negatively impacts their health [Acevedo-Garcia and Lochner, 2003], their mortality [Jackson *et al.*, 2000], and in general their socioeconomic conditions [Massey and Denton, 1993]. While in the early days of research on segregation the emergence of homogeneous neighborhoods was attributed to the individual intolerance of the citizens, it was shown by [Schelling, 1971] that residential segregation also emerges in a tolerant population. In his landmark model, he considers two types of agents that live on a line or a grid as residential area. Every agent has a tolerance level τ and is content with her position, if at least a τ-fraction of her direct neighbors are of her type. Discontent agents randomly jump to other empty positions or swap positions with another discontent agent. Schelling found that even for $\tau < \frac{1}{2}$, i.e., even if everyone is content with being in the minority within her neighborhood, random initial placements are over time transformed to placements having large homogeneous regions, i.e., many agents are surrounded by same-type neighbors, by the individual random movements of the agents. It is important to note that the agent behavior is driven by a slight bias towards preferring a certain number of same-type neighbors and that this bias on the microlevel is enough to tip the macrolevel state towards segregation. Schelling coined the term “micromotives versus macrobehavior” for such phenomena [Schelling, 2006].

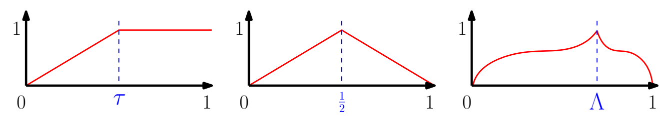

Figure 1: The x-axis shows the fraction of same-type neighbors, the y-axis the utility. Left: example of the monotone utility functions employed in recent related work. Middle and right: example of a single-peaked utility function considered in this paper.

Since its inception, Schelling’s influential model was thoroughly studied by sociologists, mathematicians and physicists via computer simulations. But only in the last decade progress has been made to understand the involved random process from a theoretical point of view. Even more recently, the Algorithmic Game Theory and the AI communities became interested in residential segregation and game-theoretic variants of Schelling’s model were studied [Chauhan *et al.*, 2018; Echzell *et al.*, 2019; Bilò *et al.*, 2020; Elkind *et al.*, 2021; Kanellopoulos *et al.*, 2021; Bullinger *et al.*, 2021]. In these strategic games the agents do not perform random moves but rather jump or swap to positions that maximize their utility. These models incorporate utility functions that are monotone in the fraction of same-type neighbors, i.e., the utility of an agent is proportional to the fraction of same-type neighbors in her neighborhood. See Figure 1 (left). However,

[1]See the racial dot map [Cai, 2013] at https://demographics.coopercenter.org/racial-dot-map/ for examples from the US.

representative sociological polls, in particular data from the General Social Survey[2] (GSS) [Smith *et al.*, 2019], indicate that this assumption of monotone utility functions should be challenged. For example, in 1982 all black respondents where asked *"If you could find the housing that you would want and like, would you rather live in a neighborhood that is all black; mostly black; half black, half white; or mostly white?"* and 54% responded with *"half black, half white"* while only 14% chose *"all black"*. Later, starting from 1988 until 2018 all respondents (of whom on average 78% were white) where asked what they think of *"Living in a neighborhood where half of your neighbors were blacks?"* a clear majority[3] responded *"strongly favor"*, *"favor"* or *"neither favor nor oppose"*. This shows that the maximum utility should not be attained in a homogeneous neighborhood.

Based on these real-world empirical observations, this paper sets out to explore a game-theoretic variant of Schelling's model with non-monotone utility functions. In particular, we will focus on single-peaked utility functions with maximum utility at a Λ-fraction of same-type neighbors (see Figure 1 (middle and right)), with $\Lambda \in (0, 1)$, satisfying mild assumptions. More precisely, we only require a function $p(x)$ to be zero-valued at $x = 0, 1$, to be strictly increasing in the interval $[0, \Lambda]$ and to be such that $p(x) = p(\frac{\Lambda(1-x)}{1-\Lambda})$ for each $x \in [\Lambda, 1]$, that is, both sides of p approach the peak, one from the left and the other from the right, in the same way, up to a rescaling due to the width of their domains ($[0, \Lambda]$, vs. $[\Lambda, 1]$). Our main findings shed light on the existence of equilibrium states and their quality in terms of the recently defined degree of integration [Elkind *et al.*, 2021] that measures the number of agents that live in a heterogeneous neighborhood.

1.1 Model

We consider a strategic game played on a given underlying connected graph $G = (V, E)$, with $|V| = n$ and $|E| = m$. For any node $v \in V$, let the *closed neighborhood* of v in G be $N[v] = \{v\} \cup \{u \in V : \{v, u\} \in E\}$ where $\delta(v) = |N[v]| - 1$ denotes the *degree of* v, and $\delta(G)$ and $\Delta(G)$ denote the minimum and the maximum degree over all nodes in G, respectively. We call a graph G *δ-regular*, if $\delta(G) = \Delta(G) = \delta$, and *almost regular*, if $\Delta(G) - \delta(G) \le 1$. We denote with $\alpha(G)$ the *independence number* of G, i.e., the cardinality of the maximum independent set in G.

A *Single-Peaked Swap Schelling Game* (G, b, Λ), called *the game*, is defined by a graph $G = (V, E)$, a positive integer $b \le n/2$ and a peak position Λ. There are n strategic agents who choose nodes in V such that every node is occupied by exactly one agent. Every agent belongs to one of two types that are associated with the colors blue and red. There are b blue agents and $r = n - b$ red agents, with blue being the color of the minority type. Let $c(i)$ be the color of agent i.

A *strategy profile* $\boldsymbol{\sigma}$ is an n-dimensional vector where all strategies are pairwise disjoint, i.e., $\boldsymbol{\sigma}$ is a permutation of V. The i-th entry of $\boldsymbol{\sigma}$ corresponds to the strategy of the i-th agent. We treat $\boldsymbol{\sigma}$ as a bijective function mapping agents to nodes, with $\boldsymbol{\sigma}^{-1}$ being its inverse function. Thus, any strategy profile $\boldsymbol{\sigma}$ corresponds to a bi-coloring of G in which exactly b nodes of G are colored blue and $n - b$ are colored red. We say that agent i *occupies node* v *in* $\boldsymbol{\sigma}$ if the i-th entry of $\boldsymbol{\sigma}$, denoted as $\boldsymbol{\sigma}(i)$, is v and, equivalently, if $\boldsymbol{\sigma}^{-1}(v) = i$. We use the notation $1_{ij}(\boldsymbol{\sigma})$, with $1_{ij}(\boldsymbol{\sigma}) = 1$ if agents i and j occupy adjacent nodes in $\boldsymbol{\sigma}$ and $1_{ij}(\boldsymbol{\sigma}) = 0$ otherwise. When $1_{ij}(\boldsymbol{\sigma}) = 1$, we say that agents are *adjacent*.

For an agent i and a feasible strategy profile $\boldsymbol{\sigma}$, we denote the set of nodes of G which are occupied by agents having the same color as agent i by $C_i(\boldsymbol{\sigma}) = \{v \in V : c(\boldsymbol{\sigma}^{-1}(v)) = c(i)\}$. Observe that $C_i(\boldsymbol{\sigma})$ includes node $\boldsymbol{\sigma}(i)$. Let $f_i(\boldsymbol{\sigma}) := \frac{|N[\boldsymbol{\sigma}(i)] \cap C_i(\boldsymbol{\sigma})|}{|N[\boldsymbol{\sigma}(i)]|}$ be the fraction of agents of her own color in i's neighborhood including herself. Thus, agents are aware of their own contribution to the diversity of their neighborhood The utility of an agent i in $\boldsymbol{\sigma}$ is defined as $\mathsf{U}_i(\boldsymbol{\sigma}) = p(f_i(\boldsymbol{\sigma}))$, where p is a single-peaked function with domain $[0, 1]$ and peak at $\Lambda \in (0, 1)$ that satisfies the following two properties: (i) p is a strictly monotonically increasing function in the interval $[0, \Lambda]$ with $p(0) = 0$; (ii) for each $x \in [\Lambda, 1]$, $p(x) = p(\frac{\Lambda(1-x)}{1-\Lambda})$ and $p(\Lambda) = 1$. Each agent aims at maximizing her utility. We say an agent i is *below the peak* when $f_i(\boldsymbol{\sigma}) < \Lambda$, *above the peak* when $f_i(\boldsymbol{\sigma}) > \Lambda$, at *the peak* when $f_i(\boldsymbol{\sigma}) = \Lambda$, and *segregated* when $f_i(\boldsymbol{\sigma}) = 1$. A game (G, b, Λ) depends also on the choice of p. However, as all our results are independent of p, we remove it from the notation for the sake of simplicity.

An agent can change her strategy only via a *swap*, i.e., exchanging node occupation with another agent. Consider two agents i and j, on nodes $\boldsymbol{\sigma}(i)$ and $\boldsymbol{\sigma}(j)$, respectively, performing a swap. This yields the new strategy profile $\boldsymbol{\sigma}_{ij}$. As agents are strategic, we only consider *profitable swaps*, i.e., swaps which strictly increase the utility of both agents. Hence, profitable swaps can only occur between agents of different colors. A strategy profile $\boldsymbol{\sigma}$ is a *swap equilibrium* (SE), if $\boldsymbol{\sigma}$ does not admit profitable swaps, i.e., if for each pair of agents i, j, we have $\mathsf{U}_i(\boldsymbol{\sigma}) \ge \mathsf{U}_i(\boldsymbol{\sigma}_{ij})$ or $\mathsf{U}_j(\boldsymbol{\sigma}) \ge \mathsf{U}_j(\boldsymbol{\sigma}_{ij})$.

We measure the quality of a strategy profile $\boldsymbol{\sigma}$ via the *degree of integration* (DoI), defined by the number of non-segregated agents. The DoI is a simple segregation measure that captures how many agents have contact with other-type agents. We prefer it to the standard utilitarian welfare since it measures segregation independently of the value of Λ. For any fixed game (G, b, Λ) let $\boldsymbol{\sigma}^*$ denote a feasible strategy profile maximizing the DoI and let $\mathsf{SE}(G, b, \Lambda)$ denote the set of swap equilibria for (G, b, Λ). We study the impact of the agents' selfishness by analyzing the *Price of Anarchy* (PoA), which is defined as $\mathsf{PoA}(G, b, \Lambda) = \frac{\mathsf{DoI}(\boldsymbol{\sigma}^*)}{\min_{\boldsymbol{\sigma} \in \mathsf{SE}(G,b,\Lambda)} \mathsf{DoI}(\boldsymbol{\sigma})}$ and the *Price of Stability* (PoS), which is defined as $\mathsf{PoS}(G, b, \Lambda) = \frac{\mathsf{DoI}(\boldsymbol{\sigma}^*)}{\max_{\boldsymbol{\sigma} \in \mathsf{SE}(G,b,\Lambda)} \mathsf{DoI}(\boldsymbol{\sigma})}$.

We investigate the *finite improvement property (FIP)* [Monderer and Shapley, 1996], i.e., if every sequence of profitable swaps is finite, which is equivalent to the existence of an ordinal potential function. For this, let $\Phi(\boldsymbol{\sigma}) = \left|\left\{\{u, v\} \in E : c(\boldsymbol{\sigma}^{-1}(u)) = c(\boldsymbol{\sigma}^{-1}(v))\right\}\right|$, counting the number of *monochromatic edges* of G under $\boldsymbol{\sigma}$,

[2]Since 50 years the GSS is regularly conducted in the US and it is a valuable and widely used data set for social scientists.

[3]In numbers: 1988: 57%, 1998: 70%, 2008: 79%, 2018: 82%. In 2018 33% answered with "favor" or "strongly favor".

i.e., the edges whose endpoints are occupied by agents of the same color, the *potential function of* $\boldsymbol{\sigma}$.

1.2 Related Work

In the last decade progress has been made to thoroughly understand the involved random process in Schelling's influential model, e.g., [Brandt *et al.*, 2012; Barmpalias *et al.*, 2014; Immorlica *et al.*, 2017].

[Zhang, 2004a; 2004b] investigated the random Schelling process via evolutionary game theory. In particular, [Zhang, 2004b] proposes a model that is similar to our model. There, agents on a toroidal grid graph with degree 4 also have a non-monotone single-peaked utility function. However, in contrast to our model, random noise is added to the utilities and transferable utilities are assumed. Zhang analyzes the Markov process of random swaps and shows that this process converges with high probability to segregated states.

The investigation of game-theoretic models for residential segregation was initiated by [Chauhan *et al.*, 2018]. There, agents are equipped with a utility function as shown in Figure 1 (left) and the finite improvement property and the PoA in terms of the number of content agents is studied. Later, [Echzell *et al.*, 2019] significantly extended these results and generalized them to games with more than two agent types.

[Elkind *et al.*, 2021] introduce a simplified model with $\tau = 1$. They prove results on the existence of equilibria, in particular that equlibria are not guaranteed to exist on trees, and on the complexity of deciding equilibrium existence. Moreover, they study the PoA in terms of the utilitarian social welfare and in terms of the newly introduced degree of integration, that counts the number of non-segregated agents. For the latter, they give a tight bound of $\frac{n}{2}$ on the PoA and the PoS that is achieved on a tree. In contrast, they derive a constant PoS on paths. [Bilò *et al.*, 2020] strengthened the PoA results for the simplified swap version w.r.t. the utilitarian social welfare function and investigated the model on almost regular graphs, grids and paths. Additionally, they introduce a variant with locality.

Recently, a model was introduced where the agent itself is included in the computation of the fraction of same-type neighbors [Kanellopoulos *et al.*, 2021]. We adopt this modified version also in our model. [Bullinger *et al.*, 2021] consider the number of agents with non-zero utility as social welfare function. They prove hardness results for computing the social optimal state and they discuss other stability notations such as Pareto optimality.

Also related are hedonic games [Drèze and Greenberg, 1980; Bogomolnaia and Jackson, 2002] where selfish agents form coalitions and the utility of an agent only depends on her coalition. Especially close are hedonic diversity games [Bredereck *et al.*, 2019; Boehmer and Elkind, 2020], where agents of different types form coalitions and the utility depends also on the type distribution in a coalition.

Our main focus is on single-peaked utility functions. This can be understood as single-peaked preferences, which date back to [Black, 1948] and are a common theme in the Economics and Game Theory literature. In particular, such preferences yield favorable behavior in the above mentioned hedonic diversity games and in the realm of voting and social choice [Walsh, 2007; Yu *et al.*, 2013; Betzler *et al.*, 2013; Elkind *et al.*, 2014; Brandt *et al.*, 2015].

1.3 Our Contribution

In this work we initiate the study of game-theoretic models for residential segregation with non-monotone utility functions. This departs from the recent line of work focusing on monotone utility functions and it opens up a promising research direction. Non-monotone utility functions are well-justified by real-world data and hence might be more suitable for modeling real-world segregation.

We focus on a broad class of non-monotone utility functions well-known in Economics and Algorithmic Game Theory: single-peaked utilities. We emphasize that our results hold for all such functions that satisfy our mild assumptions. See Table 1 for a detailed result overview.

For games with integration-oriented agents, i.e., $\Lambda \leq 1/2$, we show that swap equilibria exist on almost regular graphs and that improving response dynamics are guaranteed to converge to such stable states. Moreover, for $\Lambda = \frac{1}{2}$ swap equilibria exist on the broad class of graphs that admit an independent set that is large enough to accommodate the minority type agents. In particular, this implies equilibrium existence and efficient computability on bipartite graphs, including trees, which is in contrast to the non-existence results by [Elkind *et al.*, 2021].

Another contrast are our bounds on the PoA. On general graphs we prove a tight bound on the PoA that depends on b, the number of agents of the minority color, and we give a bound of $\Delta(G)$ on all graphs G, that is asymptotically tight on δ-regular graphs. Also for the PoS we get stronger positive results compared to [Elkind *et al.*, 2021]. For $\Lambda = \frac{1}{2}$ we give a tight PoS bound of 2 on bipartite graphs and show that the PoS is 1 on almost regular graphs with maximum degree 3, or if the size of the maximum independent set of the graph is at most b. The latter implies a PoS of 1 on regular graphs for balanced games, i.e., if there are equally many agents of both colors. Even more general, for constant $\Lambda \leq \frac{1}{2}$ we prove a constant PoS on almost regular graphs via a sophisticated proof technique that relies on the greedy algorithm for the k-MAX-CUT problem.

Additional complexity results and all omitted details can be found in [Bilò *et al.*, 2022].

2 Preliminaries

In this section, we provide some facts and lemmas that will be widely exploited throughout the paper. We start by observing the following fundamental relationship occurring between $f_i(\boldsymbol{\sigma})$ and $f_j(\boldsymbol{\sigma}_{ij})$ for two swapping agents i and j:

$$\text{if } f_i(\boldsymbol{\sigma}) = \frac{x}{y},^4 \text{ then } f_j(\boldsymbol{\sigma}_{ij}) = \frac{y + 1 - x - 1_{ij}(\boldsymbol{\sigma})}{y}. \quad (1)$$

Using property (1), we claim the following observation.

[4] For the sake of conciseness, from now on, whenever we write $f_i(\boldsymbol{\sigma}) = x/y$ for some agent i, we implicitly mean that $x := |N[\boldsymbol{\sigma}(i)] \cap C_i(\boldsymbol{\sigma})|$ and $y := |N[\boldsymbol{\sigma}(i)]|$. Observe that, under this assumption, $f_i(\boldsymbol{\sigma}) = 3/6$ is different than $f_i(\boldsymbol{\sigma}) = 1/2$.

graph classes	Equilibrium Existence		Finite Improvement Property	
arbitrary	$\times$ (Thm. 2)	$\Lambda > 1/2$	$\times$ (Thm. 2, 3)	$\Lambda \geq 1/2$
	✓(Thm. 4)	$\frac{1}{\delta(G)+1} \leq \Lambda \leq 1/2$, $\alpha(G) + 1 \geq b$		
bipartite	✓(Cor. 1)	$\Lambda = 1/2$		
1-**regular**	✓(Thm. 1)	$\Lambda \leq 1/2$	✓(Thm. 1)	$\Lambda \leq 1/2$
2-**regular**	$\times$ (Thm. 2)	$\Lambda > 1/2$	$\times$ (Thm. 2)	$\Lambda > 1/2$
	Price of Anarchy		**Price of Stability**	
arbitrary	$\leq \min\{\Delta(G), \frac{n}{b+1}, \frac{(\Delta+1)b}{b+1}\}$ (Thm. 5)		$\geq \Omega\left(\sqrt{n\Lambda}\right)$ (Thm. 9)	
bipartite	$\geq \frac{n-1}{3}$ (Thm. 5)	$b = 1$	2 (Thm. 10, 11)	$\Lambda = 1/2$
	$\geq \frac{n}{b+1}$ (Thm. 5)	$b > 1$		
regular	$\leq \min\{(\delta+1)/2, n/2b\}$ (Thm. 6)	$\Lambda < 1/\delta$		
	$\geq \frac{\delta+1}{2} - \frac{\delta+1}{4\delta+2}$ (Thm. 7)	$\Lambda \leq 1/2, \delta \geq 2$		
1-regular			1 (Thm. 12), 13)	$\Lambda \leq 1/2, \Delta(G) \leq 3$ or $\Lambda \in \left[\frac{1}{\delta(G)+1}, 1/2\right], b \geq \alpha(G)$
			$\min\{\Delta(G)+1, \mathrm{O}(1/\Lambda)\}$ (Thm. 14)	$\Lambda \leq 1/2, b < \alpha(G)$
			$\mathrm{O}(1)$ (Cor. 3)	$\Lambda \leq 1/2$
ring	$> 2 - \epsilon$ (Thm. 8)			
	$> 3/2 - \epsilon$ (Thm. 8)	$\Lambda < 1/2$		

Table 1: Result overview. We investigate the existence of equilibria, the finite improvement property, the PoA and the PoS. The "✓" symbol denotes that the respective property holds, the "$\times$" means the opposite. The respective conditions are stated next to the result. ϵ is a constant larger than zero. "1-regular" stands for almost regular graphs. Note, PoS results for almost regular graphs hold for regular graphs as well. For the PoA the stated lower bounds of other graph classes hold for arbitrary graphs as well.

Observation 1. *If $f_i(\boldsymbol{\sigma}) = x/y < 1/2$, then $f_j(\boldsymbol{\sigma}_{ij}) > 1/2$. If $f_i(\boldsymbol{\sigma}) = x/y > 1/2$, then $f_j(\boldsymbol{\sigma}_{ij}) \leq 1/2$ unless $y = 2x-1$ and $1_{ij}(\boldsymbol{\sigma}) = 0$, for which $f_j(\boldsymbol{\sigma}_{ij}) = f_i(\boldsymbol{\sigma}) = x/y > 1/2$.*

The following series of lemmas characterizes the conditions under which a profitable swap can take place.

Lemma 1. *For any $\Lambda \leq 1/2$, no profitable swaps can occur between agents below the peak.*

Lemma 2. *For any $\Lambda \leq 1/2$, no profitable swaps can occur between adjacent agents at different sides of the peak.*

Proof. Assume towards a contradiction, that i and j can perform a profitable swap in $\boldsymbol{\sigma}$, and, w.l.o.g., that $f_i(\boldsymbol{\sigma}) = x/y < \Lambda$ and $f_j(\boldsymbol{\sigma}) = x'/y' > \Lambda$. By Observation 1, j ends up above the peak in $\boldsymbol{\sigma}_{ij}$. As j improves after the swap, we have $\mathsf{U}_j(\boldsymbol{\sigma}_{ij}) = p(1 - x/y) > \mathsf{U}_j(\boldsymbol{\sigma}) = p(x'/y')$ which, given that $1 - x/y > \Lambda$ and $x'/y' > \Lambda$, yields $1 - x/y < x'/y'$. This implies that $f_i(\boldsymbol{\sigma}_{ij}) = 1 - x'/y' < 1 - 1 + x/y = x/y = f_i(\boldsymbol{\sigma})$ which, given that $f_i(\boldsymbol{\sigma}) < \Lambda$, contradicts the fact that i improves after the swap. □

Lemma 3. *For any $\Lambda \leq 1/2$, no profitable swaps can occur between agents at different sides of the peak in games on almost regular graphs.*

3 Existence of Equilibria

In this section, we provide existential results for games played on some specific graph topologies. We start by showing that games on almost regular graphs enjoy the FIP property and converge to a SE in at most m swaps in any game in which the peak does not exceed $1/2$. This result does not hold when the peak exceeds $1/2$, as we prove the existence of a game played on a 2-regular graph (i.e., a ring) admitting no SE.

Theorem 1. *For any $\Lambda \leq 1/2$, fix a game (G, b, Λ) on an almost regular graph G and a strategy profile $\boldsymbol{\sigma}$. Any sequence of profitable swaps starting from $\boldsymbol{\sigma}$ ends in a SE after at most m swaps.*

Proof. We show that, after a profitable swap, Φ decreases by at least 1. Consider a profitable swap performed by agents i and j such that $f_i(\boldsymbol{\sigma}) = x/y$ and $f_j(\boldsymbol{\sigma}) = x'/(y+t)$, with $t \in \{0, 1\}$ since G is almost regular. By Lemmas 1 and 3, we have that both i and j are above the peak, i.e., $x/y > \Lambda$ and $x'/(y+t) > \Lambda$. Thus, a necessary condition for the swap to be profitable is that $f_i(\boldsymbol{\sigma}_{ij}) < f_i(\boldsymbol{\sigma})$ and $f_j(\boldsymbol{\sigma}_{ij}) < f_j(\boldsymbol{\sigma})$. By Observation 1, the latter yields $x'/(y+t) > 1 - x/y + (1 - 1_{ij}(\boldsymbol{\sigma}))/y$, which gives $x' > y - x + 1 - 1_{ij}(\boldsymbol{\sigma}) + t(1 - x/y + (1 - 1_{ij}(\boldsymbol{\sigma}))/y) \geq y - x + 1 - 1_{ij}(\boldsymbol{\sigma})$. Since x, x', y and $1_{ij}(\boldsymbol{\sigma})$ are integers, we derive $x' \geq y - x + 2 - 1_{ij}(\boldsymbol{\sigma})$. As it holds that $\Phi(\boldsymbol{\sigma}) - \Phi(\boldsymbol{\sigma}_{ij})$ equals $x - 1 + x' - 1 - (y - x - 1_{ij}(\boldsymbol{\sigma}) + y + t - x' - 1_{ij}(\boldsymbol{\sigma})) = 2(x + x' - 1 + 1_{ij}(\boldsymbol{\sigma})) - 2y - t$, we get $\Phi(\boldsymbol{\sigma}) - \Phi(\boldsymbol{\sigma}_{ij}) \geq 1$. □

Theorem 2. *For any $\Lambda > 1/2$, there exists a game played on a 2-regular graph admitting no SE.*

Proof. Consider an instance of a game played on a ring with 6 nodes, where $b = r = 3$. Only the following two complementary cases may occur: Either, the blue agents occupy nodes that induce a path of length 2. In this case, there are two segregated agents of different colors, both with utility 0. As $p(0) = 0$ and $p(x) > 0$ for $x \in (0, 1)$, the two agents swap their positions. Or, there are two neighboring agents i and j of different colors being below the peak. In this case, as $p(1/3) < p(2/3)$, both i and j prefer to swap their positions. □

A fundamental question is whether a SE always exists in games with tolerant agents, i.e., for $\Lambda \leq 1/2$. Next result shows that Theorem 1 cannot be generalized to all graphs.

Theorem 3. *There cannot exist an ordinal potential function in games on arbitrary graphs for* $\Lambda = 1/2$*.*

For the special case of $\Lambda = 1/2$, however, existence of a SE is guaranteed in any graph whose independence number is at least the number of blue agents.

Theorem 4. *Fix a game* (G, b, Λ) *with* $\frac{1}{\delta(G)+1} \leq \Lambda \leq 1/2$*. Any strategy profile in which all agents of a same color are located on an independent set of G is a SE.*

Proof. Let $\boldsymbol{\sigma}$ be a strategy profile in which all agents of a same color are located on an independent set of G. Assume, w.l.o.g., that all blue agents are assigned to the nodes of an independent set of G and consider a profitable swap performed by a blue agent i and a red agent j. If $1_{ij}(\boldsymbol{\sigma}) = 0$, since i is only adjacent to red agents other than j, it holds that $f_j(\boldsymbol{\sigma}_{ij}) = 1$, which gives $\mathsf{U}_j(\boldsymbol{\sigma}_{ij}) = 0$, thus contradicting the fact that j performs a profitable swap. If $1_{ij}(\boldsymbol{\sigma}) = 1$, instead, we obtain $f_i(\boldsymbol{\sigma}) = \frac{1}{\delta(\boldsymbol{\sigma}(i))+1} \leq \frac{1}{\delta(G)+1} \leq \Lambda$. The numerator comes from the fact that i is only adjacent to red agents. Knowing that i cannot be at the peak, we conclude that she is below the peak. If j is also below the peak, Lemma 1 contradicts the fact that the swap is profitable, while, if j is above the peak, the contradiction comes from Lemma 2. □

Corollary 1. *For* $\Lambda = 1/2$*, games played on bipartite graphs always admit a SE which can be efficiently computed.*

4 Price of Anarchy

In this section, we give bounds on the PoA for games played on different topologies, even in those cases for which existence of a SE in not guaranteed.

4.1 General Graphs

Next lemmas provide a necessary condition that needs to be satisfied by any SE and an upper bound of the value on the social optimum, respectively.

Lemma 4. *In a SE for any game* (G, b, Λ)*, no agents of different colors can be segregated.*

Proof. Fix a strategy profile $\boldsymbol{\sigma}$. If there exist two agents i and j such that $f_i(\boldsymbol{\sigma}) = f_j(\boldsymbol{\sigma}) = 1$, they can perform a profitable swap, as $f_i(\boldsymbol{\sigma}) = f_j(\boldsymbol{\sigma}) = 1$ and $f_i(\boldsymbol{\sigma}_{ij}) = f_j(\boldsymbol{\sigma}_{ij}) \notin \{0, 1\}$. So, $\boldsymbol{\sigma}$ cannot be a SE for (G, b, Λ). □

(a) An instance with $b = 1$ blue agents. Left: $\boldsymbol{\sigma}^*$ with $\mathsf{DoI}(\boldsymbol{\sigma}^*) = n - 1$. Right: a SE $\boldsymbol{\sigma}$ with $\mathsf{DoI}(\boldsymbol{\sigma}) = 3$.

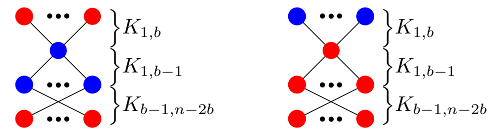

(b) An instance with $b \geq 2$ blue agents. Left: $\boldsymbol{\sigma}^*$ with $\mathsf{DoI}(\boldsymbol{\sigma}^*) = n$. Right: a SE $\boldsymbol{\sigma}$ with $\mathsf{DoI}(\boldsymbol{\sigma}) = b + 1$.

Figure 2: Lower bounds for $\mathsf{PoA}(G, b, \Lambda)$ when (a) $b = 1$, and (b) $b > 1$. Left: the socially optimal placement $\boldsymbol{\sigma}^*$. Right: the SE $\boldsymbol{\sigma}$ with minimum social welfare.

Lemma 5. *For any game* (G, b, Λ)*, we have* $\mathsf{DoI}(\boldsymbol{\sigma}^*) \leq \min\{(\Delta(G) + 1)b, n\}$*.*

Proof. As a blue node can be adjacent to at most $\Delta(G)$ red ones, it follows that, in any strategy profile, there cannot be more than $(\Delta(G) + 1)b$ non-segregated agents, so that $\mathsf{DoI}(\boldsymbol{\sigma}^*) \leq \min\{(\Delta(G) + 1)b, n\}$. □

We now give (almost) tight bounds on the PoA.

Theorem 5. *For any game* (G, b, Λ)*,* $\mathsf{PoA}(G, b, \Lambda) \leq \min\{\Delta(G), \frac{n}{b+1}, \frac{(\Delta+1)b}{b+1}\}$*. Moreover, there exists a game on a bipartite graph such that* $\mathsf{PoA}(G, b, \Lambda) \geq \frac{n}{b+1}$ *when* $b > 1$ *and* $\mathsf{PoA}(G, b, \Lambda) \geq \frac{n-1}{3}$ *when* $b = 1$*.*

Proof. For the upper bound, fix a game (G, b, Λ) and a SE $\boldsymbol{\sigma}$. By Lemma 4, only agents of one color, say c, can be segregated in $\boldsymbol{\sigma}$. Thus, we get $\mathsf{DoI}(\boldsymbol{\sigma}) \geq b + 1$. Let V be the set of nodes of color $c' \neq c$. Every node in V has to be adjacent to a node of color c. So, there are at least $|V| \geq b$ non-monochromatic edges in the coloring induced by $\boldsymbol{\sigma}$. As every node of color c can be adjacent to at most $\Delta(G)$ nodes of color c', there must be at least $\lceil b/\Delta(G) \rceil$ nodes of color c incident to a non-monochromatic edge, that is, being non-segregated in $\boldsymbol{\sigma}$. Thus, we get $\mathsf{DoI}(\boldsymbol{\sigma}) \geq \frac{(\Delta(G)+1)b}{\Delta(G)}$. We conclude that $\mathsf{DoI}(\boldsymbol{\sigma}) \geq \max\{\frac{(\Delta(G)+1)b}{\Delta(G)}, b + 1\}$. The upper bounds follow from Lemma 5. For the lower bounds, consider the games defined in Figure 2. □

4.2 Regular Graphs

For δ-regular graphs, we derive an upper bound of δ on the PoA from Theorem 5. A better result is possible when Λ is sufficiently small.

Theorem 6. *For any game* (G, b, Λ) *on a* δ*-regular graph* G *with* $\Lambda < 1/\delta$*,* $\mathsf{PoA}(G, b, \Lambda) \leq \min\{(\delta + 1)/2, n/2b\}$*.*

As a lower bound, we have the following.

Theorem 7. *For every $\delta \geq 2$ and $\Lambda \leq 1/2$, there exists a game (G, b, Λ) on a δ-regular graph such that* $\mathsf{PoA}(G, b, \Lambda) \geq \frac{\delta(\delta+1)}{2\delta+1} = \frac{\delta+1}{2} - \frac{\delta+1}{4\delta+2}$.

The lower bound given in Theorem 7 holds for all values of δ. It may be the case then that, for fixed values of δ, better bounds are possible. For $\delta = 2$ indeed, lower bounds matching the upper bounds given in Theorems 5 and 6 can be derived.

Theorem 8. *For any $\epsilon > 0$, there exists a game (G, b, Λ) on a ring such that* $\mathsf{PoA}(G, b, 1/2) > 2 - \epsilon$ *and* $\mathsf{PoA}(G, b, \Lambda) > 3/2 - \epsilon$ *for* $\Lambda < 1/2$.

5 Price of Stability

In this section, we give bounds on the PoS for games played on different topologies.

5.1 General Graphs

We give a lower bound on the PoS on general graphs which asymptotically matches the upper bound on the PoA when $b = \Theta(\sqrt{n})$ and Λ is a constant w.r.t n.

Theorem 9. *For every Λ, there is a game (G, b, Λ) such that* $\mathsf{PoS}(G, b) = \Omega(\sqrt{n\Lambda})$.

5.2 Bipartite Graphs

For bipartite graphs, we provide a tight bound of 2 for the PoS of games for which the peak is at $1/2$. We start with the upper bound.

Theorem 10. *For any game $(G, b, 1/2)$ on a bipartite graph G, we have* $\mathsf{PoS}(G, b, 1/2) \leq 2$.

We now give the matching lower bound.

Theorem 11. *There exists a game $(G, b, 1/2)$ on a bipartite graph such that* $\mathsf{PoS}(G, b, 1/2) \geq 2$.

5.3 Almost Regular Graphs

We provide upper bounds to the PoS for games played on almost regular graphs. We start by considering the case of graphs with small degree.

Theorem 12. *For any game (G, b, Λ) on an almost regular graph with $\Delta(G) \leq 3$ and $\Lambda \leq 1/2$,* $\mathsf{PoS}(G, b, \Lambda) = 1$.

An analogous result holds for the case in which $b \geq \alpha(G)$.

Theorem 13. *For any game (G, b, Λ) on an almost regular graph with $b \geq \alpha(G)$ and $\frac{1}{\delta(G)+1} \leq \Lambda \leq 1/2$, we have* $\mathsf{PoS}(G, b, \Lambda) = 1$.

A game (G, b, Λ) is *balanced* if $b = \lfloor n/2 \rfloor$. Using Theorem 13, we show that the PoS is 1 in balanced games on regular graphs.

Corollary 2. *For any balanced game (G, b, Λ) on a δ-regular graph G and $\frac{1}{\delta+1} \leq \Lambda \leq 1/2$, we have* $\mathsf{PoS}(G, b, \Lambda) = 1$.

Proof. We have that $b = \lfloor n/2 \rfloor$. We show that $\alpha(G) \leq \lfloor n/2 \rfloor$ using a simple counting argument. This allows us to use Theorem 13 to prove the claim.

To show the upper bound on $\alpha(G)$, we count all the edges that are incident to the nodes of a fixed maximum independent set of G and bound this value from above by the number of edges of the graph, thus obtaining the following inequality $\delta\alpha(G) \leq \frac{\delta}{2}n$, i.e., $\alpha(G) \leq n/2$. Using the fact that $\alpha(G)$ is an integer value, we derive $\alpha(G) \leq \lfloor n/2 \rfloor$. □

We now give the upper bound to the PoS for games played on almost regular graphs when $b < \alpha(G)$.

Theorem 14. *For any game (G, b, Λ) on an almost regular graph G with $b < \alpha(G)$ and $\Lambda \leq 1/2$, we have* $\mathsf{PoS}(G, b, \Lambda) = \min\{\Delta(G) + 1, O(1/\Lambda)\}$.

We can derive the following upper bound to the PoS.

Corollary 3. *For any game (G, b, Λ) on an almost regular graph with a constant value of $\Lambda \leq 1/2$, we have* $\mathsf{PoS}(G, b, \Lambda) = O(1)$.

Proof. By Theorem 5, the PoS is constant if $\Delta(G)$ is constant. The result when $\Delta(G)$ is not constant is divided into two cases. For the case $b \geq \alpha(G)$ the claim immediately follows from Theorem 13. For the case $b < \alpha(G)$ the claim follows from Theorem 14 and the fact that Λ is constant by assumption. □

6 Conclusion and Future Work

We study game-theoretic residential segregation with integration-oriented agents and thereby open up the novel research direction of considering non-monotone utility functions. Our results clearly show that moving from monotone to non-monotone utilities yields novel structural properties and different results in terms of equilibrium existence and quality. We have equilibrium existence for a larger class of graphs, compared to [Elkind *et al.*, 2021], and it is an important open problem to prove or disprove if swap equilibria for our model with $\Lambda \leq \frac{1}{2}$ are guaranteed to exist on any graph.

So far we considered single-peaked utilities that are supported by data from real-world sociological polls. However, also other natural types of non-monotone utilities could be studied. Also ties in the utility function could be resolved by breaking them consistently towards favoring being in the minority or being in the majority. The non-existence example of swap equilibria used in the proof of Theorem 2 also applies to the case with $\Lambda = \frac{1}{2}$ and breaking ties towards being in the majority. Interestingly, by breaking ties the other way we get the same existence results as without tie-breaking and also our other results hold in this case. This is another indication that tolerance helps with stability.

Moreover, all our existence results also hold for utility functions having a symmetric plateau shape around Λ. Investigating the PoA for these utility functions is open.

Regarding the quality of the equilibria, we analyze the degree of integration as social welfare function, as this is inline with considering integration-oriented agents. Of course, studying the quality of the equilibria in terms of the standard utilitarian social welfare, i.e., $\mathsf{SUM}(\boldsymbol{\sigma}) = \sum_{i=1}^{n} \mathsf{U}_i(\boldsymbol{\sigma})$, would also be interesting. We note in passing that on ring topologies the PoA and the PoS with respect to both social welfare functions coincide.

References

[Acevedo-Garcia and Lochner, 2003] Dolores Acevedo-Garcia and Kimberly A. Lochner. Residential segregation and health. *Neighborhoods and Health*, pages 265–87, 2003.

[Barmpalias *et al.*, 2014] George Barmpalias, Richard Elwes, and Andy Lewis-Pye. Digital morphogenesis via schelling segregation. In *FOCS 2014*, pages 156–165, 2014.

[Betzler *et al.*, 2013] Nadja Betzler, Arkadi Slinko, and Johannes Uhlmann. On the computation of fully proportional representation. *JAIR*, 47:475–519, 2013.

[Bilò *et al.*, 2020] Davide Bilò, Vittorio Bilò, Pascal Lenzner, and Louise Molitor. Topological influence and locality in swap schelling games. In *MFCS 2020*, pages 15:1–15:15, 2020.

[Bilò *et al.*, 2022] Davide Bilò, Vittorio Bilò, Pascal Lenzner, and Louise Molitor. Tolerance is necessary for stability: Single-peaked swap schelling games. *CoRR*, abs/2204.12599, 2022.

[Black, 1948] Duncan Black. On the rationale of group decision-making. *J. Pol. E.*, 56(1):23–34, 1948.

[Boehmer and Elkind, 2020] Niclas Boehmer and Edith Elkind. Individual-based stability in hedonic diversity games. In *AAAI 2020*, pages 1822–1829, 2020.

[Bogomolnaia and Jackson, 2002] Anna Bogomolnaia and Matthew O. Jackson. The stability of hedonic coalition structures. *GEB*, 38(2):201–230, 2002.

[Brandt *et al.*, 2012] Christina Brandt, Nicole Immorlica, Gautam Kamath, and Robert Kleinberg. An analysis of one-dimensional schelling segregation. In *STOC 2012*, pages 789–804, 2012.

[Brandt *et al.*, 2015] Felix Brandt, Markus Brill, Edith Hemaspaandra, and Lane A. Hemaspaandra. Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. *JAIR*, 53:439–496, 2015.

[Bredereck *et al.*, 2019] Robert Bredereck, Edith Elkind, and Ayumi Igarashi. Hedonic diversity games. In *AAMAS 2019*, pages 565–573, 2019.

[Bullinger *et al.*, 2021] Martin Bullinger, Warut Suksompong, and Alexandros A. Voudouris. Welfare guarantees in schelling segregation. *J. Artif. Intell. Res.*, 71:143–174, 2021.

[Cai, 2013] Qian Cai. The racial dot map. *Weldon Cooper Center for Public Service, University of Virginia*, 2013.

[Chauhan *et al.*, 2018] Ankit Chauhan, Pascal Lenzner, and Louise Molitor. Schelling segregation with strategic agents. In *SAGT 2018*, pages 137–149, 2018.

[Drèze and Greenberg, 1980] Jaques H. Drèze and Jay Greenberg. Hedonic coalitions: Optimality and stability. *Econometrica: J. Eco. Soc.*, pages 987–1003, 1980.

[Echzell *et al.*, 2019] Hagen Echzell, Tobias Friedrich, Pascal Lenzner, Louise Molitor, Marcus Pappik, Friedrich Schöne, Fabian Sommer, and David Stangl. Convergence and hardness of strategic schelling segregation. In *WINE 2019*, pages 156–170, 2019.

[Elkind *et al.*, 2014] Edith Elkind, Piotr Faliszewski, and Piotr Skowron. A characterization of the single-peaked single-crossing domain. In *AAAI 2014*, 2014.

[Elkind *et al.*, 2021] Edith Elkind, Jiarui Gan, Ayumi Igarashi, Warut Suksompong, and Alexandros A. Voudouris. Schelling games on graphs. *Artif. Intell.*, 301:103576, 2021.

[Immorlica *et al.*, 2017] Nicole Immorlica, Robert Kleinberg, Brendan Lucier, and Morteza Zadomighaddam. Exponential segregation in a two-dimensional schelling model with tolerant individuals. In *SODA 2017*, pages 984–993, 2017.

[Jackson *et al.*, 2000] Sharon A. Jackson, Roger T. Anderson, Norman J. Johnson, and Paul D. Sorlie. The relation of residential segregation to all-cause mortality: A study in black and white. *Am. J. Pub. Health*, 90(4):615, 2000.

[Kanellopoulos *et al.*, 2021] Panagiotis Kanellopoulos, Maria Kyropoulou, and Alexandros A. Voudouris. Modified schelling games. *TCS*, 880:1–19, 2021.

[Massey and Denton, 1988] Douglas S. Massey and Nancy A. Denton. The dimensions of residential segregation. *Soc. Forces*, 67(2):281–315, 1988.

[Massey and Denton, 1993] Douglas S. Massey and Nancy A. Denton. *American Apartheid: Segregation and the Making of the Underclass*. Harvard U. Press, 1993.

[Monderer and Shapley, 1996] Dov Monderer and Lloyd S. Shapley. Potential games. *GEB*, 14(1):124–143, 1996.

[Schelling, 1971] Thomas C. Schelling. Dynamic models of segregation. *J. Math. Soc.*, 1(2):143–186, 1971.

[Schelling, 2006] Thomas C. Schelling. *Micromotives and Macrobehavior*. WW Norton & Company, 2006.

[Smith *et al.*, 2019] Tom W. Smith, Michael Davern, Jeremy Freese, and Stephen L. Morgan. General social surveys, 1972–2018 cumulative codebook. *NORC ed. Chicago: NORC 2019, U. Chicago*, 2019.

[Walsh, 2007] Toby Walsh. Uncertainty in preference elicitation and aggregation. In *AAAI 2007*, pages 3–8, 2007.

[Yu *et al.*, 2013] Lan Yu, Hau Chan, and Edith Elkind. Multiwinner elections under preferences that are single-peaked on a tree. In *AAAI 2013*, pages 425–431, 2013.

[Zhang, 2004a] Junfu Zhang. A dynamic model of residential segregation. *J. Mat. Soc.*, 28(3):147–170, 2004.

[Zhang, 2004b] Junfu Zhang. Residential segregation in an all-integrationist world. *J. Econ. Behav. and Org.*, 54(4):533–550, 2004.