

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

Davide Bilò¹, Vittorio Bilò², Pascal Lenzner³ and Louise Molitor³

¹University of L'Aquila, Italy, davide.bilo@univaq.it

²University of Salento, Italy, vittorio.bilo@unisalento.it

³Hasso Plattner Institute, University of Potsdam, Germany, {firstname.lastname}@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 time¹.

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

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

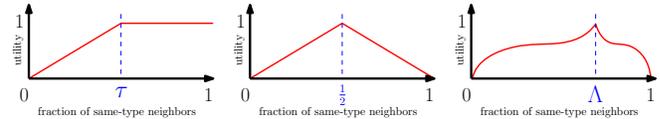


Figure 1: Left: example of the monotone utility functions employed in recent related work. Middle and right: example of a single-peaked utility function considered in this paper.

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].

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; Agarwal *et al.*, 2021; Kanellopoulos *et al.*, 2021a; 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, representative sociological polls, in particular data from the

General Social Survey² (GSS) [Smith *et al.*, 2019], indicate that this assumption of monotone utility functions should be challenged. For example, in 1982 all black respondents were 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) were asked what they think of “Living in a neighborhood where half of your neighbors were blacks?” a clear majority³ 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 [Agarwal *et al.*, 2021] that measures the number of agents that live in a heterogeneous neighborhood.

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 $\Delta(G) - \delta(G) \leq 1$. 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 $b \leq n/2$ 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* σ is an n -dimensional vector where all strategies are pairwise disjoint, i.e., σ is a permutation of V .

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

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

The i -th entry of σ corresponds to the strategy of the i -th agent. We treat σ as a bijective function mapping agents to nodes, with σ^{-1} being its inverse function. Thus, any strategy profile σ 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 σ if the i -th entry of σ , denoted as $\sigma(i)$, is v and, equivalently, if $\sigma^{-1}(v) = i$. We use the notation $1_{ij}(\sigma)$, with $1_{ij}(\sigma) = 1$ if agents i and j occupy adjacent nodes in σ and $1_{ij}(\sigma) = 0$ otherwise. When $1_{ij}(\sigma) = 1$, we say that agents are *adjacent*.

For an agent i and a feasible strategy profile σ , we denote the set of nodes of G which are occupied by agents having the same color as agent i by

$$C_i(\sigma) = \{v \in V : c(\sigma^{-1}(v)) = c(i)\}.$$

Observe that $C_i(\sigma)$ includes node $\sigma(i)$. Let

$$f_i(\sigma) := \frac{|N[\sigma(i)] \cap C_i(\sigma)|}{|N[\sigma(i)]|}$$

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 σ is defined as $U_i(\sigma) = p(f_i(\sigma))$, 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 $f_i(\sigma) < \Lambda$, *above the peak* when $f_i(\sigma) > \Lambda$, *at the peak* when $f_i(\sigma) = \Lambda$, and *segregated* when $f_i(\sigma) = 1$. 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 $\sigma(i)$ and $\sigma(j)$, respectively, performing a swap. This yields the new strategy profile σ_{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 σ is a *swap equilibrium* (SE), if σ does not admit profitable swaps, i.e., if for each pair of agents i, j , we have $U_i(\sigma) \geq U_i(\sigma_{ij})$ or $U_j(\sigma) \geq U_j(\sigma_{ij})$.

We measure the quality of a strategy profile σ via the *degree of integration* (Dol), defined by the number of non-segregated agents. The Dol 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 σ^* denote a feasible strategy profile maximizing the Dol and let $SE(G, b, \Lambda)$ 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

$$\text{PoA}(G, b, \Lambda) = \frac{\text{Dol}(\sigma^*)}{\min_{\sigma \in SE(G, b, \Lambda)} \text{Dol}(\sigma)}$$

and the *Price of Stability* (PoS), which is defined as

$$\text{PoS}(G, b, \Lambda) = \frac{\text{Dol}(\sigma^*)}{\max_{\sigma \in \text{SE}(G, b, \Lambda)} \text{Dol}(\sigma)}.$$

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

$$\Phi(\sigma) = |\{\{u, v\} \in E : c(\sigma^{-1}(u)) = c(\sigma^{-1}(v))\}|,$$

counting the number of *monochromatic edges* of G under σ , i.e., the edges whose endpoints are occupied by agents of the same color, the *potential function* of σ .

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; Barmalinas *et al.*, 2014; Immorlica *et al.*, 2017].

Zhang [2004a,b] 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.

Agarwal *et al.* [2021] introduce a simplified model with $\tau = 1$. They prove results on the existence of equilibria, in particular that equilibria 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. The complexity results were extended by Kreisel *et al.* [2021]. Chan *et al.* [2020] and Kanellopoulos *et al.* [2021b] considered generalized variants.

Recently, a model was introduced where the agent itself is included in the computation of the fraction of same-type neighbors [Kanellopoulos *et al.*, 2021a]. 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].

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 Agarwal *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 [Agarwal *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.

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

graph classes	Equilibrium Existence		Finite Improvement Property	
arbitrary	× (Thm. 2)	$\Lambda > 1/2$	× (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	× (Thm. 2)	$\Lambda > 1/2$	× (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)		$\geq \Omega(\sqrt{n\Lambda})$ (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 $\Lambda \in [\frac{1}{\delta(G)+1}, 1/2], b \geq \alpha(G)$
			$\min\{\Delta(G) + 1, O(1/\Lambda)\}$ (Thm. 14)	$\Lambda \leq 1/2, b < \alpha(G)$
			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 “×” 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.

$f_i(\sigma)$ and $f_j(\sigma_{ij})$ for two swapping agents i and j :

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

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

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

Proof. If $f_i(\sigma) = x/y < 1/2$, we have

$$f_j(\sigma_{ij}) = \frac{y+1-x-1_{ij}(\sigma)}{y} = 1 - \frac{x}{y} + \frac{1-1_{ij}(\sigma)}{y} > \frac{1}{2}.$$

If $f_i(\sigma) = x/y > 1/2$, we distinguish among different cases. If $1_{ij}(\sigma) = 1$, we get

$$f_j(\sigma_{ij}) = \frac{y+1-x-1_{ij}(\sigma)}{y} = 1 - \frac{x}{y} < \frac{1}{2},$$

while, if $y < 2x-1$, which implies $x \geq \frac{y+2}{2}$, it follows

$$\begin{aligned} f_j(\sigma_{ij}) &= \frac{y+1-x-1_{ij}(\sigma)}{y} \\ &= 1 - \frac{x}{y} + \frac{1-1_{ij}(\sigma)}{y} \\ &\leq \frac{1}{2} - \frac{1_{ij}(\sigma)}{y} \leq \frac{1}{2}. \end{aligned}$$

Finally, for $y = 2x-1$ and $1_{ij}(\sigma) = 0$, we get

$$f_j(\sigma_{ij}) = \frac{y+1-x-1_{ij}(\sigma)}{y} = \frac{x}{y} > \frac{1}{2}.$$

□

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

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.*

Proof. Fix a strategy profile σ and two agents i and j , below the peak, who can perform a profitable swap in σ . By Observation 1, both i and j are above the peak in σ_{ij} . Assume, w.l.o.g., that $f_i(\sigma) = x/y < \Lambda$ which yields

$$f_j(\sigma_{ij}) = \frac{y+1-x-1_{ij}(\sigma)}{y} > \Lambda.$$

We claim that $U_j(\sigma_{ij}) \leq U_i(\sigma)$. By the definition of p , this holds whenever

$$\begin{aligned} \frac{x}{y} &\geq \frac{\Lambda}{1-\Lambda} \left(1 - \frac{y+1-x-1_{ij}(\sigma)}{y} \right) \\ &= \frac{\Lambda}{1-\Lambda} \left(\frac{x}{y} - \frac{1-1_{ij}(\sigma)}{y} \right) \end{aligned}$$

which holds true as $\frac{\Lambda}{1-\Lambda} \leq 1$ and $1-1_{ij}(\sigma) \geq 0$.

By applying the same argument, with i and j swapped, we also get $U_i(\sigma_{ij}) \leq U_j(\sigma)$. As the swap is profitable, we have $U_i(\sigma) < U_i(\sigma_{ij})$ and $U_j(\sigma) < U_j(\sigma_{ij})$. Putting all these inequalities together, we conclude that

$$U_j(\sigma_{ij}) \leq U_i(\sigma) < U_i(\sigma_{ij}) \leq U_j(\sigma) < U_j(\sigma_{ij}),$$

which yields a contradiction. □

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 σ , and, w.l.o.g., that $f_i(\sigma) =$

$x/y < \Lambda$ and $f_j(\sigma) = x'/y' > \Lambda$. By Observation 1, j ends up above the peak in σ_{ij} . As j improves after the swap, we have

$$U_j(\sigma_{ij}) = p(1 - x/y) > U_j(\sigma) = p(x'/y')$$

which, given that $1 - x/y > \Lambda$ and $x'/y' > \Lambda$, yields

$$1 - x/y < x'/y'.$$

This implies that

$$f_i(\sigma_{ij}) = 1 - x'/y' < 1 - 1 + x/y = x/y = f_i(\sigma)$$

which, given that $f_i(\sigma) < \Lambda$, contradicts the fact that i improves after the swap. \square

In the following we present a technical result that will be helpful in proving Lemma 4.

Lemma 3. *For any $\Lambda \leq 1/2$, any profitable swap occurring between two agents i and j in a strategy profile σ , with $f_i(\sigma) < \Lambda$ and $f_j(\sigma) > \Lambda$, requires $\delta(\sigma(i)) > \delta(\sigma(j))$.*

Proof. Assume towards a contradiction, that

$$\delta(\sigma(i)) \leq \delta(\sigma(j))$$

and i and j can perform a profitable swap in σ , and, w.l.o.g., that $f_i(\sigma) = x/y < \Lambda$ and $f_j(\sigma) = x'/y' > \Lambda$.

By Lemma 2, it must be $1_{ij}(\sigma) = 0$. By Observation 1, j ends up above the peak in σ_{ij} . As j improves after the swap, we have

$$U_j(\sigma_{ij}) = p(1 - x/y + 1/y) > U_j(\sigma) = p(x'/y')$$

which, given that $1 - x/y + 1/y > \Lambda$ and $x'/y' > \Lambda$, yields

$$1 - x/y + 1/y < x'/y'.$$

This implies that

$$\begin{aligned} f_i(\sigma_{ij}) &= 1 - x'/y' + 1/y' \\ &< 1 - 1 + x/y - 1/y + 1/y' \\ &= x/y + 1/y' - 1/y. \end{aligned}$$

Now, as the hypothesis $\delta(\sigma(i)) \leq \delta(\sigma(j))$ can be restated as $y \leq y'$, we derive

$$f_i(\sigma_{ij}) < x/y + 1/y' - 1/y \leq x/y = f_i(\sigma),$$

which, given that $f_i(\sigma) < \Lambda$, contradicts the fact that i improves after the swap. \square

Lemma 4. *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.*

Proof. Fix a strategy profile σ and two agents i and j at different sides of the peak admitting a profitable swap in σ . As the game is played on an almost regular graph, by Lemma 3, it must be $f_i(\sigma) = x/(y+t) < \Lambda$, $f_j(\sigma) = x'/y > \Lambda$ and $t \in \{0, 1\}$. Moreover, by Lemma 2, we have $1_{ij}(\sigma) = 0$.

Since j improves after the swap, we have

$$U_j(\sigma_{ij}) = p((y+t-x+1)/(y+t)) > U_j(\sigma) = p(x'/y)$$

which, given that $(y+t-x+1)/(y+t) > \Lambda$ and $x'/y > \Lambda$, yields

$$x'/y > (y+t-x+1)/(y+t).$$

We derive $x'(y+t) > y(y+t-x+1)$, which, given that both sides of the inequality are integers, yields

$$x'(y+t) \geq y(y+t-x+1) + 1. \quad (2)$$

Since i improves after the swap, we have

$$U_i(\sigma_{ij}) = p((y-x'+1)/y) > U_i(\sigma) = p(x/(y+t)).$$

We now distinguish between two possible cases:

(i) $(y-x'+1)/y \leq \Lambda$ and

(ii) $(y-x'+1)/y > \Lambda$.

If case (i) occurs, it must be

$$x/(y+t) < (y-x'+1)/y$$

which is equivalent to

$$x'(y+t) < (y+t)(y+1) - xy.$$

Together with inequality (2), this yields

$$y(y+t-x+1) + 1 < (y+t)(y+1) - xy$$

which is satisfied if and only if $t > 1$. Given that $t \in \{0, 1\}$, we derive a contradiction.

If case (ii) occurs, from $(y-x'+1)/y > \Lambda$, we get

$$x' < y(1-\Lambda) + 1. \quad (3)$$

Since $U_i(\sigma_{ij}) = p((y-x'+1)/y) > U_i(\sigma) = p(x/(y+t))$, $x/(y+t) < \Lambda$ and $(y-x'+1)/y > \Lambda$, by the definition of p , we derive

$$x/(y+t) < \frac{\Lambda}{1-\Lambda} (1 - (y-x'+1)/y) = \frac{\Lambda}{1-\Lambda} (x'-1)/y,$$

by which we get $x < \frac{\Lambda(x'-1)(y+t)}{(1-\Lambda)y}$. Together with inequality (3), this yields

$$x < \Lambda(y+t). \quad (4)$$

By summing up inequalities (3) and (4), we get

$$x + x' < y + 1 + \Lambda t.$$

As x, x', y are integers, $t \in \{0, 1\}$ and $\Lambda \in [0, 1/2]$, we derive

$$x + x' \leq y + 1. \quad (5)$$

Starting from inequality (2) and then using inequality (5), we derive

$$\begin{aligned} x't &\geq y(y+t-x-x'+1) + 1 \\ &\geq yt + 1, \end{aligned}$$

which, given that $x' \leq y$, is never satisfied when $t \in \{0, 1\}$. Thus, also in this case, we derive a contradiction. \square

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 σ . Any sequence of profitable swaps starting from σ 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 $f_i(\sigma) = x/y$ and $f_j(\sigma) = x'/(y+t)$, with $t \in \{0, 1\}$ since G is almost regular. By Lemmas 1 and 4, 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 $f_i(\sigma_{ij}) < f_i(\sigma)$ and $f_j(\sigma_{ij}) < f_j(\sigma)$. By Observation 1, the latter yields

$$x'/(y+t) > 1 - x/y + (1 - 1_{ij}(\sigma))/y,$$

which gives

$$\begin{aligned} x' &> y - x + 1 - 1_{ij}(\sigma) + t(1 - x/y + (1 - 1_{ij}(\sigma))/y) \\ &\geq y - x + 1 - 1_{ij}(\sigma). \end{aligned}$$

Since x, x', y and $1_{ij}(\sigma)$ are integers, we derive

$$x' \geq y - x + 2 - 1_{ij}(\sigma).$$

As it holds that $\Phi(\sigma) - \Phi(\sigma_{ij})$ equals

$$\begin{aligned} &x - 1 + x' - 1 - \\ &(y - x - 1_{ij}(\sigma) + y + t - x' - 1_{ij}(\sigma)) \\ &= 2(x + x' - 1 + 1_{ij}(\sigma)) - 2y - t, \end{aligned}$$

we get $\Phi(\sigma) - \Phi(\sigma_{ij}) \geq 1$. \square

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. \square

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 $\Lambda = 1/2$.*

Proof. We prove the statement by providing an improving response cycle where in every step a profitable swap is possible. The construction and arising strategy profiles are shown in Figure 2. For the sake of simplicity we assume $p(x) = x$ for $x \in [0, 1/2]$. However, our result is independent of the choice of p . Remember that, since $\Lambda = \frac{1}{2}$, $p(x) = p(1-x)$ for $x \in [1/2, 1]$.

In the initial placement, cf. Figure 2a, agents a and b can swap. By swapping their positions, agent a can increase her utility from $1 - \frac{5}{7} = \frac{2}{7}$ to $1 - \frac{9}{13} = \frac{4}{13}$ and agent b increases her utility from $\frac{5}{13}$ to $\frac{3}{7}$.

Next, agents c and d can swap, cf. Figure 2b. Swapping with agent d increases agent c 's utility from $1 - \frac{4}{7} = \frac{3}{7}$ to $\frac{6}{13}$, and agent d can increase her utility from $1 - \frac{8}{13} = \frac{5}{13}$ to $1 - \frac{4}{7} = \frac{3}{7}$.

After this, cf. Figure 2c agents e and f , and agents g and h , respectively, have the opportunity to swap and increase their utility. Agent e increases her utility from $\frac{3}{7}$ to $\frac{1}{2}$ while agent f increases her utility from $1 - \frac{2}{2} = 0$ to $1 - \frac{5}{7} = \frac{2}{7}$. Agent g improves her utility from $1 - \frac{2}{2} = 0$ to $1 - \frac{3}{5} = \frac{2}{5}$ and agent h increases her utility from $1 - \frac{3}{5} = \frac{2}{5}$ to $\frac{1}{2}$.

Next, cf. Figure 2d, swaps between the agents i and j , and k and l , respectively, are possible. Agent i can increase her utility from $1 - \frac{2}{2} = 0$ to $1 - \frac{4}{5} = \frac{1}{5}$ and j can increase her utility from $\frac{2}{5}$ to $\frac{1}{2}$. Agent k improves her utility from $1 - \frac{2}{2} = 0$ to $1 - \frac{3}{5} = \frac{2}{5}$ and agent l increases her utility from $1 - \frac{3}{5} = \frac{2}{5}$ to $\frac{1}{2}$.

In the next step, agents e and m can swap. Agent e increases her utility from $1 - \frac{2}{2} = 0$ to $1 - \frac{3}{5} = \frac{2}{5}$ and m increases her utility from $1 - \frac{3}{5} = \frac{2}{5}$ to $\frac{1}{2}$.

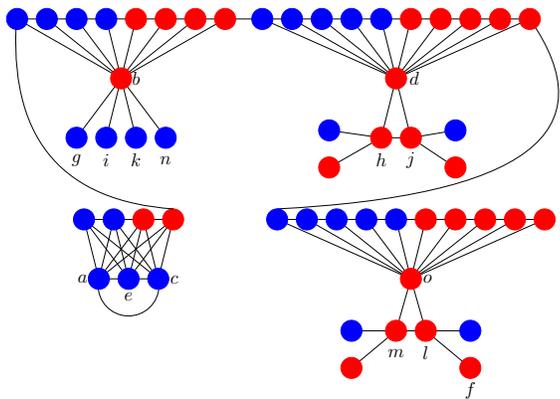
In a final step agents n and o swap. Swapping with agent o increases agent n 's utility from $1 - \frac{2}{2} = 0$ to $1 - \frac{8}{13} = \frac{5}{13}$ and agent o increases her utility from $\frac{6}{13}$ to $\frac{1}{2}$. The now reached placement, cf. Figure 2g is equivalent to the initial placement, cf. Figure 2a.

However, note that although convergence is not guaranteed there still exists a stable state, cf. Figure 2h. \square

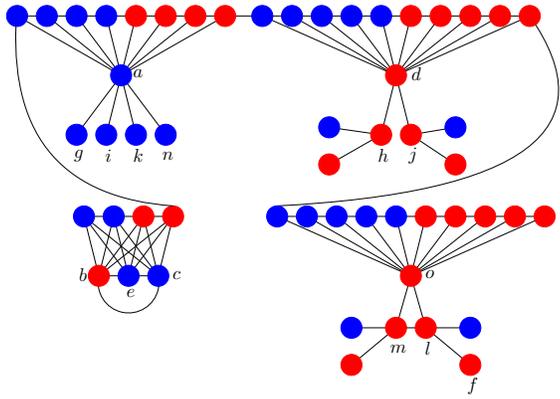
For the special case of $\Lambda = 1/2$, however, existence of a SE is guaranteed in any graph whose independence number is at least the number of blue agents.

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.*

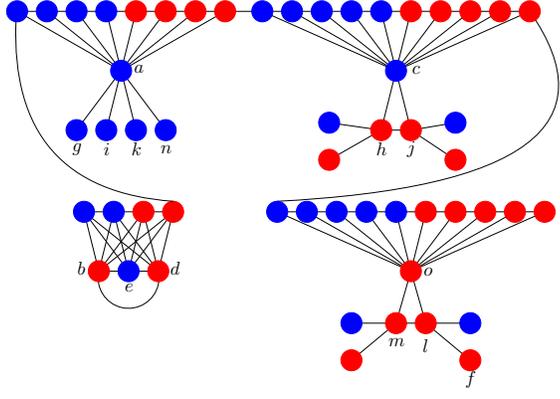
Proof. Let σ be a strategy profile in which all agents of a same color are located on an independent set of G . Assume, w.l.o.g., that all blue agents are assigned to the nodes of an independent set of G and consider a profitable swap performed by a blue agent i and a red agent j . If $1_{ij}(\sigma) = 0$, since i is only adjacent to red agents other than j , it holds that $f_j(\sigma_{ij}) = 1$, which gives $U_j(\sigma_{ij}) = 0$, thus contradicting the fact that j performs a profitable swap. If $1_{ij}(\sigma) = 1$, instead, we obtain $f_i(\sigma) = \frac{1}{\delta(\sigma(i))+1} \leq \frac{1}{\delta(G)+1} \leq \Lambda$. 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



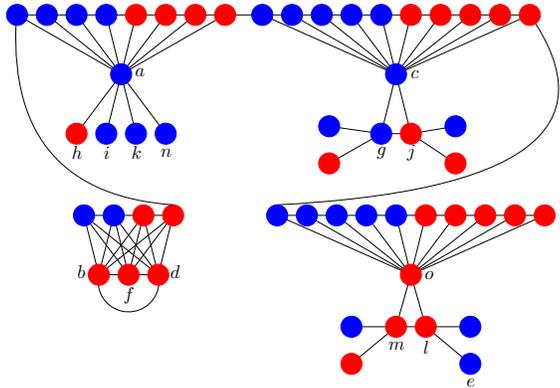
(a) Initial placement



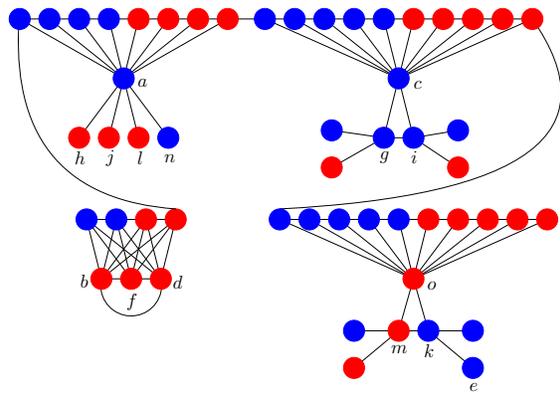
(b) Placement after the first swap of a and b



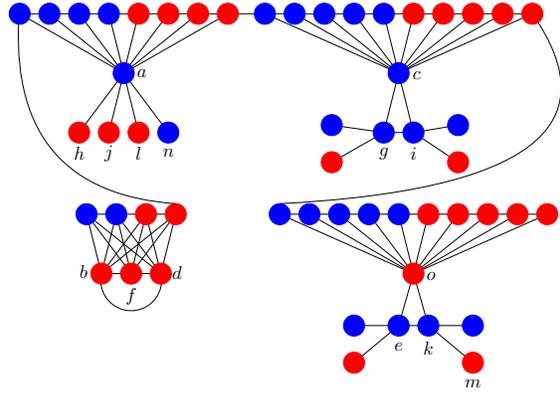
(c) Placement after the second swap of c and d



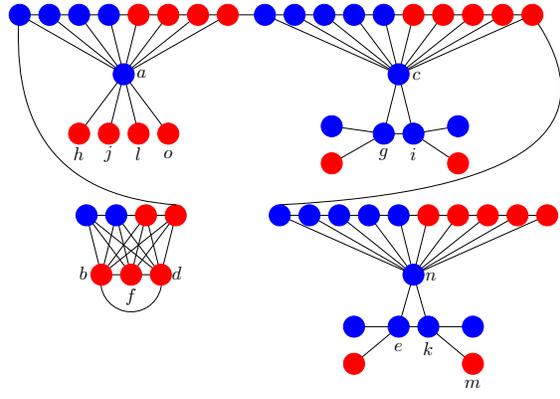
(d) Placement after the third and fourth swap of e and f , and g and h , respectively.



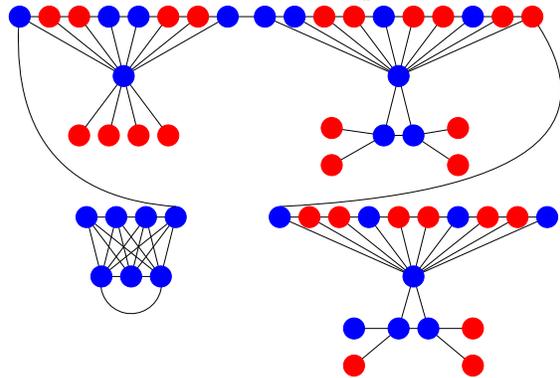
(e) Placement after the fifth and sixth swap of i and j , and k and l , respectively.



(f) Placement after the seventh swap of e and m .



(g) Placement after the last swap of n and o .



(h) The SE for the same instance.

Figure 2: An IRC and the SE for a game $(G, b, \frac{1}{2})$.

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. \square

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 is not guaranteed.

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

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

Lemma 6. *For any game (G, b, Λ) , we have*

$$\text{Dol}(\sigma^*) \leq \min\{(\Delta(G) + 1)b, n\}.$$

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 $\text{Dol}(\sigma^*) \leq \min\{(\Delta(G) + 1)b, n\}$. \square

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

Theorem 5. *For any game (G, b, Λ) ,*

$$\text{PoA}(G, b, \Lambda) \leq \min\left\{\Delta(G), \frac{n}{b+1}, \frac{(\Delta+1)b}{b+1}\right\}.$$

Moreover, there exists a game on a bipartite graph such that $\text{PoA}(G, b, \Lambda) \geq \frac{n}{b+1}$ when $b > 1$ and $\text{PoA}(G, b, \Lambda) \geq \frac{n-1}{3}$ when $b = 1$.

Proof. For the upper bound, fix a game (G, b, Λ) and a SE σ . By Lemma 5, only agents of one color, say c , can be segregated in σ . Thus, we get $\text{Dol}(\sigma) \geq b + 1$. 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 σ . 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 σ . Thus, we get

$$\text{Dol}(\sigma) \geq \frac{(\Delta(G) + 1)b}{\Delta(G)}.$$

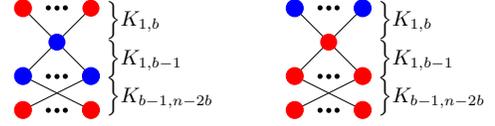
We conclude that

$$\text{Dol}(\sigma) \geq \max\left\{\frac{(\Delta(G) + 1)b}{\Delta(G)}, b + 1\right\}.$$

The upper bounds follow from Lemma 6. For the lower bounds, consider the games defined in Figure 3. \square



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



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

Figure 3: Lower bounds for $\text{PoA}(G, b, \Lambda)$ when (a) $b = 1$, and (b) $b > 1$. Left: the socially optimal placement σ^* . Right: the SE σ with minimum social welfare.

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 $\Lambda < 1/\delta$, $\text{PoA}(G, b, \Lambda) \leq \min\{(\delta + 1)/2, n/2b\}$.*

Proof. Fix a SE σ . By Lemma 5, only agents of a unique color, say c , can be segregated in σ . Let V be the set of nodes of color $c' \neq c$. As $\Lambda < 1/\delta$, every node in V has to be adjacent to nodes of color c only. Otherwise, any agent in V that is adjacent to an agent of color c' can perform a profitable swap with a segregated agent of color c . Thus, there are at least $\delta|V| \geq \delta b$ non-monochromatic edges in the coloring induced by σ . As every agent of one color can be adjacent to at most δ agents of the other one, there are at least b non-segregated agents of color c . Together with the at least b agents of color c' , this gives $\text{Dol}(\sigma) \geq 2b$ which, together with Lemma 6, yields the claim. \square

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*

$$\text{PoA}(G, b, \Lambda) \geq \frac{\delta(\delta + 1)}{2\delta + 1} = \frac{\delta + 1}{2} - \frac{\delta + 1}{4\delta + 2}.$$

Proof. Consider graph G shown in Figure 4. G is made of three combined gadgets that we call the *left gadget*, the *upper right gadget* and the *lower right gadget*. The left gadget is essentially the complete bipartite graph $K_{\delta, \delta}$ with a missing edge: the one connecting the last two nodes of the respective partitions. These nodes are connected to the upper right gadget and the lower right one, respectively. So each node in the left gadget has degree δ . The upper right gadget consists of a clique $K_{\delta-1}$ whose nodes are all connected to two special nodes, one on the left of $K_{\delta-1}$ and one on the right. The node on the left is the one adjacent to the node from the left gadget, while the node on the right connects the gadget with the lower right one. Thus, every node in this gadget has degree δ . Finally, the lower right gadget is any δ -regular graph with n'

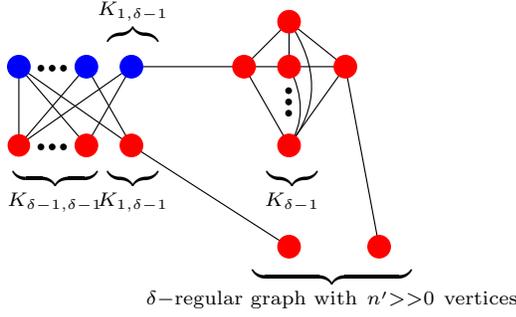


Figure 4: Lower bound construction for the PoA on δ -regular graphs. The SE σ has $\text{Dol}(\sigma) = 2\delta + 1$.

nodes with a missing edge: the one connecting the node incident to the edge coming from the left gadget with the node incident to the edge coming from the upper right gadget. So, every node in this gadget has degree equal to δ and we conclude that G is δ -regular.

Now set $b = \delta$. We claim that the strategy profile σ depicted in Figure 4 is a SE. If a red agent i swaps with a non-adjacent blue agent j , we have $f_i(\sigma_{ij}) = 1$ which results in a non-profitable swap. So a red agent can profitably swap only with an adjacent blue agent. Any blue agent j has $f_j(\sigma) = \frac{1}{\delta+1}$. By swapping with an adjacent red agent i , we have that either $f_j(\sigma_{ij}) = \frac{1}{\delta+1}$ or $f_j(\sigma_{ij}) = \frac{\delta}{\delta+1}$ which never yields an improvement when $\Lambda \leq 1/2$ or $f_j(\sigma_{ij}) = \frac{\delta-1}{\delta+1}$ when swapping one blue agent j with the rightmost red one of the left gadget. In this case, it is red agent i that has no improvement in swapping since $f_i(\sigma) = \frac{2}{\delta+1}$ and $f_i(\sigma_{ij}) = \frac{\delta}{\delta+1}$ which never yields an improvement when $\Lambda \leq 1/2$. So, σ is a SE such that $\text{Dol}(\sigma) = 2\delta + 1$.

By letting n' go to infinity, it is always possible to select $b = \delta$ nodes in the lower right gadget such that their closed neighborhoods are pairwise disjoint, which yields $\text{Dol}(\sigma^*) = \delta(\delta + 1)$ and thus the desired lower bound. \square

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 $\text{PoA}(G, b, 1/2) > 2 - \epsilon$ and $\text{PoA}(G, b, \Lambda) > 3/2 - \epsilon$ for $\Lambda < 1/2$.*

Proof. For any even $b \geq 2$, let G be a ring defined by the sequence of nodes v_1, v_2, \dots, v_n , with $n = 3b$.

Assume $\Lambda = 1/2$. In this case,

$$p(1/3) = p(2/3) > p(0) = p(1).$$

Let σ be the strategy profile obtained as follows: starting from v_1 , assign two blue agents followed by a red as long as this is possible. At this time, nodes up to v_x , with $x = 3b/2$, have been assigned an agent. All the remaining nodes are assigned the remaining red agents. Since for any blue agent i we have $f_i(\sigma) = 2/3$, all blue agents are getting the largest

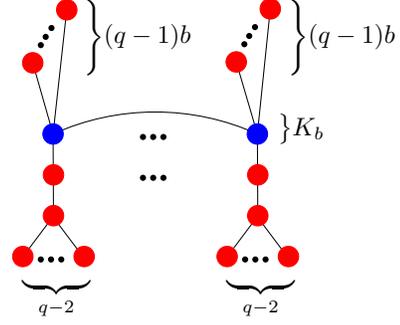


Figure 5: The instance used in the proof of Theorem 9. Shown is the socially optimal placement σ^* .

possible utility and are not interested in swapping. Thus, σ is a SE. As all agents residing at nodes v_n, v_1, \dots, v_x are not segregated, we have $\text{Dol}(\sigma) = 3b/2 + 1$.

Now assume $\Lambda < 1/2$. In this case,

$$p(1/3) > p(2/3) > p(0) = p(1).$$

Let σ be the strategy profile obtained by alternating blue and red agents for as much as possible. Since for any blue agent i we have $f_i(\sigma) = 1/3$, all blue agents are getting the largest possible utility and are not interested in swapping. Thus, σ is a SE. As all agents residing at nodes v_n, v_1, \dots, v_{2b} are not segregated, we have $\text{Dol}(\sigma) = 2b + 1$.

A strategy profile of social value $3b$ can be obtained by sequencing triplets made of two red agents with a blue one in between. Both claims follow by choosing b sufficiently large. \square

5 Price of Stability

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

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 $\text{PoS}(G, b) = \Omega(\sqrt{n\Lambda})$.*

Proof. Let $q \geq 2$ be an integer such that $\frac{1}{q} \leq \Lambda < \frac{1}{q-1}$. Consider the instance in Figure 5 in which there is a clique of b nodes such that every node in the clique is additionally adjacent to $(q-1)b$ leaves (depicted to the up) and to the leaf of a star with q nodes (depicted at the bottom). It is easy to check that $n = (q-1)b^2 + b + bq$. Letting b go to infinity, we get $n = \Theta((q-1)b^2)$, by which $b = \Omega(\sqrt{\frac{n}{q-1}}) = \Omega(\sqrt{n\Lambda})$.

The socially optimal placement σ^* , depicted in Figure 5 has all blue agents on the nodes of the clique and thereby achieves

$$\text{Dol}(\sigma^*) = (q-1)b^2 + 2b = \Omega((q-1)b^2).$$

We claim that, in contrast, any SE σ can have at most one blue agent in the clique. Assume, by way of contradiction, that there are two or more blue agents in the clique, and let i

be one of these agents. Then, there is at least one center of a star of q nodes that is occupied by a segregated red agent j . If i and j swap, we have $f_j(\sigma) = 1$ and $f_j(\sigma_{ij}) \notin \{0, 1\}$ so that j improves, and

$$f_i(\sigma) \leq \frac{b}{qb+1}$$

and

$$f_i(\sigma_{ij}) = \frac{1}{q} \in \left(\frac{b}{qb+1}, \Lambda \right],$$

so that i improves too. This contradicts that σ is a SE. Hence, for any SE σ , it holds that

$$\text{Dol}(\sigma) \leq (b-1)q + qb + 1 = O(2qb)$$

(the fact that a SE exists can be easily checked by considering the strategy profile obtained by placing one blue agent in a node v of the clique and all the remaining $b-1$ ones to the $b-1$ center of a star of q nodes that are not appended to v). It follows that the price of stability is

$$\frac{\Omega((q-1)b^2)}{O(2qb)} = \Omega(b) = \Omega(\sqrt{n\Lambda}).$$

□

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 $\text{PoS}(G, b, 1/2) \leq 2$.*

Proof. Let (V_1, V_2) , with $|V_1| \leq |V_2|$, be the bipartition of the nodes of G . For a fixed optimal profile σ^* , denote by B_1 (resp. B_2) the set of nodes of V_1 (resp. V_2) occupied by a blue agent in σ^* . Moreover, denote by R_1 (resp. R_2) the set of nodes occupied by a red agent in σ^* falling in the neighborhood of some node in B_2 (resp. B_1). Clearly, we have $\text{Dol}(\sigma^*) \leq b + |R_1| + |R_2|$. We shall prove the existence of two swap equilibria, namely σ_1 and σ_2 , whose performance compare nicely with that of σ^* .

To construct σ_1 , start from σ^* and swap all blue agents in B_2 with red agents in V_1 as long as this is possible. If all blue agents end up in V_1 , we have that all blue agents occupy the nodes of an independent set of G and so, by Theorem 4, σ_1 is an SE. If some blue agents are left out from V_1 , then all red agents are located in V_2 . So, we have that all red agents occupy the nodes of an independent set of G and, by Theorem 4, σ_1 is an SE also in this case. As the set of nodes in B_1 are blue in both σ^* and σ_1 , we obtain that $\text{Dol}(\sigma_1) \geq |R_2|$.

Equilibrium σ_2 is obtained symmetrically by swapping all blue agents in B_1 with red agents in V_2 as long as this is possible. In this case, as $b \leq n/2 \leq |V_2|$, all blue agents end up in V_2 and, by Theorem 4, σ_2 is an SE. As the set of nodes in B_2 are blue in both σ^* and σ_2 and all blue agents are adjacent to some red agent in σ_2 , we obtain that

$$\text{Dol}(\sigma_2) \geq b + |R_1|.$$

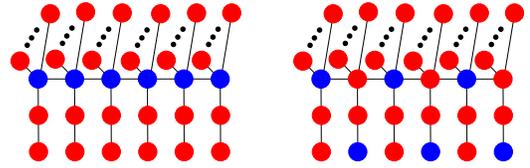


Figure 6: Left: a game with its socially optimal strategy profile σ^* shown. Right: the SE with maximum social welfare for the same instance.

Thus, we conclude that

$$\begin{aligned} \text{PoS}(G, b, 1/2) &\leq \frac{\text{Dol}(\sigma^*)}{\max\{\text{Dol}(\sigma_1), \text{Dol}(\sigma_2)\}} \\ &\leq \frac{b + |R_1| + |R_2|}{\max\{b + |R_1|, |R_2|\}} \leq 2. \end{aligned}$$

□

We now give the matching lower bound.

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

Proof. Consider the instance $(G, b, 1/2)$ defined in Figure 6. G consists of a path of b nodes, that we call the base of the graph. Any node in the base of the graph is additionally connected to $2(b-1)$ leaves (depicted on the up) and to a 2-node path (depicted on the bottom). For the socially optimal profile σ^* , we get

$$\text{Dol}(\sigma^*) = 2(b-1)b + 2b.$$

However, this is not a SE. In any strategy profile σ in which two blue agents are adjacent in the base of the graph, like in the socially optimal profile σ^* , one of them, denote this agent by i , can swap with a segregated red agent, denoted by j , placed on a leaf node in the lower row. Observe that agent j is always guaranteed to exist. Since we have

$$f_i(\sigma) \leq \frac{b}{2b+1} < 1/2$$

and

$$f_i(\sigma_{ij}) = 1/2,$$

agent i improves her utility. For agent j , we have $f_j(\sigma) = 1$ and $f_j(\sigma_{ij}) \notin \{0, 1\}$, so the swap is profitable and σ cannot be a SE. The maximum number of agents with non-zero utility that can be obtained by respecting this necessary constraint is

$$\text{Dol}(\sigma) = b(b-1) + \frac{5b}{2},$$

achieved by the SE σ depicted in Figure 6 (right). The claim follows by letting b go to infinity. □

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 $\Lambda \leq 1/2$, $\text{PoS}(G, b, \Lambda) = 1$.*

Proof. Let σ^* be a socially optimal profile. Using Lemma 8 we have that there exists a SE σ satisfying $\text{Dol}(\sigma) \geq \text{Dol}(\sigma^*)$. \square

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 $\text{PoS}(G, b, \Lambda) = 1$.*

Proof. We prove the claim by showing the existence of a SE σ such that $\text{Dol}(\sigma) = n$. Clearly, if $b = \alpha(G)$, then the strategy profile σ in which the blue agents occupy all the nodes of a maximum independent set of G is a SE by Theorem 4, and thus, $\text{Dol}(\sigma) = n$, and the statement follows. Therefore, in the following we assume that $b, r > \alpha(G)$.

Let σ be a strategy profile minimizing the value $\Phi(\sigma)$ (ties are arbitrarily broken). By Theorem 1, σ is a SE. We now prove that $\text{Dol}(\sigma) = n$. For the sake of contradiction, assume that $\text{Dol}(\sigma) < n$, i.e., there is at least a segregated agent, say i , in σ . Assume w.l.o.g. that i is blue. We claim that all red agents are placed on nodes that form an independent set, i.e., $r \leq \alpha(G)$. This allows us to obtain the desired contradiction as $r > \alpha(G)$. If the red agents are not placed on nodes that form an independent set, then there exists a red agent, say j , having at least a red neighbor in σ . The strategy profile σ_{ij} satisfies

$$\Phi(\sigma) - \Phi(\sigma_{ij}) \geq \delta(\sigma(i)) + 1 - (\delta(\sigma(j)) - 1) \geq 1,$$

since $|\delta(\sigma(i)) - \delta(\sigma(j))| \leq 1$. Therefore, $\Phi(\sigma_{ij}) < \Phi(\sigma)$, thus contradicting the fact that σ minimizes Φ . \square

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 $\text{PoS}(G, b, \Lambda) = 1$.*

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 $\alpha(G) \leq \lfloor n/2 \rfloor$. \square

Before proving that the PoS is $O(1/\Lambda)$ when $b < \alpha(G)$, we need to introduce some new definitions and additional technical lemmas based on some well-known optimization cut problems.

For a given graph G and a subset of nodes of V , we denote by $G[U]$ the subgraph of G induced by U . More precisely, the node set of $G[U]$ is U and, for every $u, v \in U$, $G[U]$ contains the edge (u, v) iff G contains the edge (u, v) .

The k -MAX-CUT problem is an optimization problem in which, given a graph G as input, we want to compute a k -partition $\{V_1, \dots, V_k\}$ of the nodes of G that maximizes the number of edges that cross the cut induced by the k -partition,

that is, the number of edges (u, v) such that $u \in V_t, v \in V_h$, and $h \neq t$. It is well-known that the greedy algorithm for the k -MAX-CUT problem computes a k -partition $\{V_1, \dots, V_k\}$ of the nodes of a graph such that, for every node v , the number of edges incident to v that cross the cut induced by the k -partition is at least $\lceil (1 - \frac{1}{k})\delta(v) \rceil$ [Vazirani, 2013]. Using this folklore result, we can derive the following useful lemma.

Observation 2. *Let G be a graph and $U \subseteq V$ such that $|U| \geq k$. There exists a k -partition $\{V_1, \dots, V_k\}$ of $G[U]$ such that, for every $t \in \{1, \dots, k\}$, the degree of each vertex $v \in V_t$ in $G[V_t]$ is at most $\lfloor \frac{\delta(v)}{k} \rfloor$.*

Proof. The greedy algorithm for the k -MAX-CUT problem computes a k -partition $\{V_1, \dots, V_k\}$ of $G[U]$ such that, for every node v of U , the number of edges incident to v that cross the cut induced by the k -partition is at least

$$\lceil (1 - \frac{1}{k})\delta(v) \rceil.$$

As a consequence, for any $v \in V_t$, with $t \in \{1, \dots, k\}$, the number of edges that are incident to v in $G[V_t]$ is at most

$$\delta(v) - \lceil (1 - \frac{1}{k})\delta(v) \rceil = \lfloor \delta(v)/k \rfloor.$$

\square

The BALANCED k -MAX-CUT problem is a k -MAX-CUT in which we additionally require the k -partition $\{V_1, \dots, V_k\}$ to be *balanced*, i.e., for every $t \in \{1, \dots, k\}$, $|V_t| \geq \lfloor n/k \rfloor$.⁵ For the BALANCED k -MAX-CUT problem we can prove a useful lemma that is analogous to Observation 2.

Lemma 7. *Let G be a graph and $U \subseteq V$ such that $|U| \geq 2$. There exists a balanced k -partition $\{V_1, \dots, V_k\}$ of U such that, for at least one $t \in \{1, \dots, k\}$, $|V_t| \geq 1$ and the degree of every $v \in V_t$ in $G[V_t]$ is at most $\lfloor \delta(v)/k \rfloor$.*

Proof. In the remainder of this proof, for a given k -partition $\{V_1, \dots, V_k\}$ of U and a bijective function

$$\rho : \{v_1, \dots, v_k\} \rightarrow \{v_1, \dots, v_k\},$$

with $v_t \in V_t$ for every $t \in \{1, \dots, k\}$, we define the ρ -swap as the k -partition $\{V'_1, \dots, V'_k\}$ in which, for every $t \in \{1, \dots, k\}$, the set V'_t is obtained from V_t by replacing v_t with $\rho(v_t)$ (it may happen that $\rho(v_t) = v_t$). We say that the ρ -swap is profitable if the number of edges crossing the cut induced by $\{V'_1, \dots, V'_k\}$ is strictly larger than the number of edges crossing the cut induced by $\{V_1, \dots, V_k\}$. A balanced k -partition $\{V_1, \dots, V_k\}$ is stable if there is no profitable ρ -swap.

Let $\{V_1, \dots, V_k\}$ be a balanced k -partition of U that maximizes the number of edges in the cut (ties can be arbitrarily broken). We claim that $\{V_1, \dots, V_k\}$ satisfies all the properties of the lemma statement. This is clearly true when $|U| \leq k$ as all edges of G are in the cut (consider, for instance, the solution in which each set of the k -partition contains at most one node). Therefore, we only need to prove the claim when $|U| > k$. This implies that $|V_t| > 0$ for every $t \in \{1, \dots, k\}$.

⁵When $n < k$, we might have empty sets in the k -partition.

As a consequence, we only need to prove that there exists $t \in \{1, \dots, k\}$ such that, for every node $v \in V_t$, the degree of v in $G[V_t]$ is at most $\lfloor \delta(v)/k \rfloor$.

For the sake of contradiction, assume that there are k nodes v_1, \dots, v_k , with $v_t \in V_t$, such that, for every $t \in \{1, \dots, k\}$, the degree of v_t in $G[V_t]$ is strictly larger than $\lfloor \delta(v_t)/k \rfloor$. Consider the graph H on the k nodes v_1, \dots, v_k , where we add the direct edge (v_t, v_h) between v_t and v_h , with $t \neq h$, iff the number of edges incident to v_t whose other endpoints are in V_h is at most $\lfloor \delta(v_t)/k \rfloor$. We observe that the outdegree of each node in H is at least 1. As a consequence H contains a directed cycle C . We define

$$\rho : \{v_1, \dots, v_k\} \rightarrow \{v_1, \dots, v_k\}$$

as follows. For each edge (v_t, v_h) of C , we define

$$\rho(v_t) = v_h,$$

while

$$\rho(v_\ell) = v_\ell$$

for every $\ell \in \{1, \dots, k\}$ such that v_ℓ is not contained in C . Clearly, ρ is a bijective function. Furthermore, the ρ -swap is profitable, as each node v_t that is moved from V_t to V'_h , with $h \neq t$, contributes with at least one more edge in the cut. This contradicts the fact that $\{V_1, \dots, V_k\}$ maximizes the number of edges in the cut. \square

We are now ready to prove the upper bound on 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 $\text{PoS}(G, b, \Lambda) = \min\{\Delta(G) + 1, O(1/\Lambda)\}$.*

Proof. Let $\Delta = \Delta(G)$. Let σ^* be a strategy profile that maximizes $\text{Dol}(\sigma^*)$ for (G, b, Λ) . Let B and R be the nodes occupied by the non-segregated blue and red agents in σ^* , respectively. Clearly, $\text{Dol}(\sigma^*) = |B| + |R|$. Moreover, $|B| \geq 1$ iff $|R| \geq 1$. We prove the claim by showing the existence of a strategy profile σ such that σ is a SE and

$$\text{Dol}(\sigma) = \Omega(\Lambda(|B| + |R|)).$$

We first rule out the case in which $b \leq \lfloor \Lambda(\Delta - 1) \rfloor - 1$. In fact, in this case $G[B]$ has maximum degree of at most $\lfloor \Lambda(\Delta - 1) \rfloor - 1$, which implies that every blue agent is below the peak in σ^* . As a consequence, from Lemma 1 and Lemma 4, σ^* is also a SE. Therefore, in the following we assume that $b > \lfloor \Lambda(\Delta - 1) \rfloor - 1$. As b is an integer, we have that $b \geq \Lambda\Delta - \Lambda - 1$, from which we derive

$$\Delta \leq b/\Lambda + 1 + 1/\Lambda. \quad (6)$$

Next, we rule out the case in which $|B| + |R| = O(b/\Lambda)$, i.e., $b = \Omega(\Lambda(|B| + |R|))$. In fact, in this case let σ be any SE (that we know to exist by Theorem 1). By Lemma 5, there is a color for which all agents of that color are not segregated in σ . Therefore

$$\text{Dol}(\sigma) \geq \min\{b, r\} = b = \omega(\Lambda(|B| + |R|)).$$

Hence, for the rest of the proof, we assume that $|B| + |R| = \omega(b/\Lambda)$. As $|B| \leq b$, it holds

$$|R| = \omega(b/\Lambda). \quad (7)$$

Finally, we rule out the case in which $\Delta = O(1/\Lambda)$. By Theorem 5, $\text{PoS}(G, b, \Lambda) \leq \Delta + 1$. In the following, for any subset B' of B , we denote by $\mathcal{R}(B')$ the subset of nodes of R that are dominated by B' , i.e., $\mathcal{R}(B')$ contains all nodes $u \in R$ for which there exists a node v in B' such that (u, v) is an edge of G .

We iteratively use Observation 2 to compute hierarchical k -partitions of B , with $k = \lceil \frac{\Delta-1}{\Lambda\Delta-1} \rceil$. We observe that $k \geq 1$ as $\Lambda > 1/(\Delta + 1)$.

Starting from B , we compute the k -partition $\{B_0^1, \dots, B_0^k\}$ of B that satisfies the premises of Observation 2. This is called the k -partition of level 0. For the remaining part of this proof we use the shortcut R_0^t to denote $\mathcal{R}(B_0^t)$. W.l.o.g., we assume that $|R_0^1| \geq \max_t |R_0^t|$. Now, given the k -partition $\{B_h^1, \dots, B_h^k\}$ of level h such that $|B_h^1| \geq 2$ and $|R_h^1| \geq \max_t |R_h^t|$, we compute a k -partition $\{B_{h+1}^1, \dots, B_{h+1}^k\}$ of B_h^1 that satisfies the premises of Observation 2. This is the k -partition of level $h+1$ where, again, we use the shortcut R_{h+1}^t to denote $\mathcal{R}(B_{h+1}^t)$ and, w.l.o.g., we assume that $|R_{h+1}^1| \geq \max_t |R_{h+1}^t|$. Clearly, the k -partition $\{B_L^1, \dots, B_L^k\}$ computed in the last level, say L , satisfies $|B_L^1| = 1$ and $|R_L^1| \leq \Delta$. We observe that

$$|R| \geq |R_0^1| \geq |R_1^1| \geq \dots \geq |R_L^1|$$

by construction.

Let ℓ be the minimum index such that $|R \setminus R_\ell^1| \geq k(b-1)$. Such an index always exists because, using inequality (7), inequality (6), and the fact that $\Delta = \omega(1/\Lambda)$, we have

$$|R \setminus R_L^1| = |R| - |R_L^1| \geq |R| - \Delta = \omega(b/\Lambda) \geq kb.$$

Let $\{R^1, \dots, R^k\}$ be a balanced k -partition of $R \setminus R_\ell^1$ that satisfies the premises of Lemma 7. We have that $|R^t| \geq b$ for every $t \in \{1, \dots, k\}$. W.l.o.g., we assume that the degree of each node $v \in R^1$ in $G[R^1]$ is at most $\lfloor \delta(v)/k \rfloor$. By construction, also each node $v \in B_\ell^1$ has a degree in $G[B_\ell^1]$ of at most $\lfloor \delta(v)/k \rfloor$. Since there is no edge (u, v) of G such that $u \in B_\ell^1$ and $v \in R \setminus R_\ell^1$, it follows that the degree of each node $v \in B_\ell^1 \cup R^1$ in $G[B_\ell^1 \cup R^1]$ is also upper bounded by $\lfloor \delta(v)/k \rfloor$.

Let σ be the strategy profile in which exactly $|B_\ell^1|$ blue agents are placed on the nodes of B_ℓ^1 and the remaining $b - |B_\ell^1|$ blue agents are placed on a subset of nodes of R^1 (ties among nodes of R^1 are arbitrarily broken). Let i be any blue agent and let v be the node occupied by i in σ . By the choice of k we have $k \geq \lceil \frac{\delta(v)}{\Lambda(\delta(v)+1)-1} \rceil$. This implies that

$$f_i(\sigma) \leq \frac{\lfloor \delta(v)/k \rfloor + 1}{\delta(v) + 1} \leq \frac{\delta(v)/k + 1}{(\delta(v) + 1)} \leq \Lambda.$$

Therefore, the blue agent i is below the peak in σ . Since every blue agent is below the peak in σ , from Lemma 1 and Corollary 4, σ is a SE.

We conclude the proof by showing that

$$\text{Dol}(\sigma) = \Omega(\Lambda(|B| + |R|)).$$

By construction, each blue agent occupies a node that is adjacent to at least one other node occupied by a red agent. Moreover, each red agent that occupies a node of R_ℓ^1 is in

the neighborhood of at least one node of B_ℓ^1 that is occupied by a blue agent). Therefore,

$$\text{Dol}(\sigma) \geq b + |R_\ell^1|. \quad (8)$$

For proof convenience, let us denote by R_{-1}^1 the set R . By the choice of ℓ we know that $|R| - |R_{\ell-1}^1| = |R \setminus R_{\ell-1}^1| < kb$, which implies that $|R_{\ell-1}^1| > |R| - kb$. As a consequence, since

$$\sum_{i=1}^k |R_\ell^i| \geq |R_{\ell-1}^1| > |R| - kb$$

and $|R_\ell^1| \geq \max_i |R_\ell^i|$, we obtain

$$|R_\ell^1| > \frac{|R|}{k} - b. \quad (9)$$

Combining the inequalities (8) and (9) we obtain

$$\text{Dol}(\sigma) > |R|/k.$$

Using inequality (7) and the fact that $k = O(1/\Lambda)$ we finally obtain $\text{Dol}(\sigma) = \Omega(\Lambda(|B| + |R|))$, as desired. \square

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 $\text{PoS}(G, b, \Lambda) = O(1)$.*

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. \square

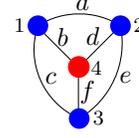
6 Computational Complexity Results

In this section we analyze the computational complexity aspects of the game played on both bipartite and regular graphs. More precisely, we provide hardness results for the two problems of computing a social optimum and a SE σ that maximizes the value $\text{Dol}(\sigma)$, respectively.

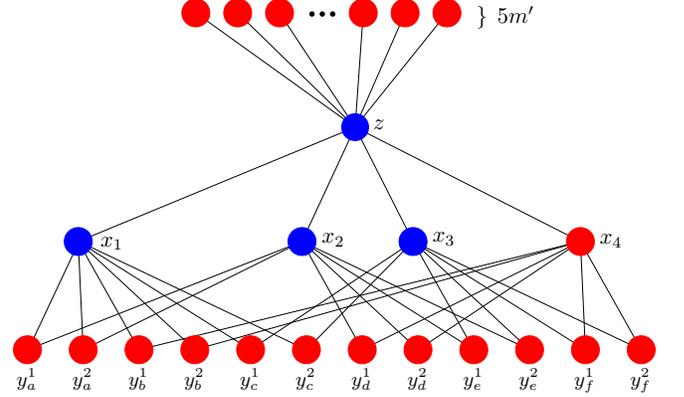
Theorem 15. *There is a constant $c > 1$ such that, given a game (G, b, Λ) played on a cubic graph G , the problem of computing a social optimum strategy profile is not c -approximable in polynomial time, unless $\text{P} = \text{NP}$.*

Proof. The reduction is from MINIMUM DOMINATING SET problem on cubic graphs, an optimization problem in which the goal is to compute a minimum-size set D of nodes of a given cubic graph G' that dominates $V(G')$, i.e., for every node $v \in V(G')$, $v \in D$ or there is an edge $(u, v) \in E(G')$ such that $u \in D$. It is known that a minimum dominating set on cubic graphs is not approximable within some constant $c' > 1$, unless $\text{P} = \text{NP}$, see [Alimonti and Kann, 1997].

Let G be a cubic graph of n nodes that has a minimum dominating set of size k^* and let $b = k^*$. We claim that a strategy profile σ^* satisfies $\text{Dol}(\sigma^*) = n$ iff the b blue agents are placed on the nodes that form a minimum dominating set of G . Indeed, a red agent placed on a node v is not segregated in σ^* iff there is a blue agent placed on a node that



(a) A cubic graph G' with $n' = 4$, $m' = 6$, with a minimum vertex cover $k^* = 3$.



(b) An instance (G, b) of our game constructed from G' with $n = n' + 7m' + 1$ vertices and $b = k^* + 1$ blue agents.

Figure 7: An example instance of the reduction from Vertex Cover shown in Theorem 16.

dominates v . Furthermore, a blue agent placed on a node u is never segregated in σ^* because of the minimality of the dominating set (i.e., each node of a minimum dominating set D must dominate a node of the graph that is not in D).

Let $c = \frac{4}{5-c'}$. Since all the nodes of the graph trivially form a dominating set of size $n \leq 4k^*$ (each node of the dominating set dominates 4 nodes), we have that $c' < 4$ and therefore, $c > 1$.

We complete the proof by showing that, if we were able to compute, in polynomial time, a strategy profile σ such that

$$\frac{\text{Dol}(\sigma^*)}{\text{Dol}(\sigma)} \leq c,$$

then we could compute, in polynomial time, a c' -approximate dominating set of G .

Let σ be a strategy profile such that $\frac{\text{Dol}(\sigma^*)}{\text{Dol}(\sigma)} \leq c$ and let D be the set of nodes that are occupied by the blue agents in σ . Let n' be the nodes of G that are not dominated by D . We have that $\frac{n}{n-n'} \leq c$, from which we derive that

$$n' \leq \frac{c-1}{c} n = \frac{\frac{4}{5-c'} - 1}{\frac{4}{5-c'}} n \leq \frac{c'-1}{4} 4k^* = (c'-1)k^*.$$

We now compute, in polynomial time, a dominating set D' of G whose size is at most $k^* + (c'-1)k^* = c'k^*$. D' contains D and all the n' nodes of G that are not dominated by D . Clearly, D' is a dominating set of G that approximates the value k^* within a factor of c' . This completes the proof. \square

The following lemma allows us to convert any strategy profile into an SE without increasing the number of segregated agents.

Lemma 8. *Given a game (G, b, Λ) on an almost regular graph, with $\Delta(G) \leq 3$ and $\Lambda \leq 1/2$, and given a strategy profile σ , we can compute an SE σ' such that $\text{Dol}(\sigma') \geq \text{Dol}(\sigma)$ in polynomial time.*

Proof. We shall prove the following claim: if a feasible strategy profile σ is not a SE, then there exists a (not necessarily profitable) swap decreasing the potential function Φ and not creating new segregated agents. This implies that after a sequence of at most $m = |E(G)|$ swaps of this type, we obtain an SE σ' such that $\text{Dol}(\sigma') \geq \text{Dol}(\sigma)$. Therefore, given σ , we have that σ' can be computed in polynomial time.

It remains to prove the existence of a (not necessarily profitable) swap σ' such that $\Phi(\sigma') < \Phi(\sigma)$ and not creating new segregated agents. Towards this end, fix a non-equilibrium feasible strategy profile σ and consider a blue agent i and a red agent j possessing a profitable swap in σ . If no segregated agents are created in σ_{ij} , then the claim holds. So assume that a segregated agent k is created in σ_{ij} . Clearly, by definition of profitable swaps, it must be $k \notin \{i, j\}$. Assume, w.l.o.g, that k is red. Then, since we have

$$f_k(\sigma_{ij}) = \frac{\delta(\sigma(k)) + 1}{\delta(\sigma(k)) + 1} = 1,$$

k needs be adjacent to i in σ , i.e. $1_{ik}(\sigma) = 1$, and

$$f_k(\sigma) = \frac{\delta(\sigma(k))}{\delta(\sigma(k)) + 1}.$$

Let x_b (resp. x_r) be the number of blue agents (resp. red agents other than k) adjacent to i in σ . Since profitable swaps in almost regular graphs can only occur between agents above the peak, we have

$$\frac{x_b + 1}{x_b + x_r + 2} > \Lambda$$

which implies $x_r < x_b$ as $\Lambda \leq 1/2$. By swapping i and k , we get

$$\Phi(\sigma) - \Phi(\sigma_{ik}) = \delta(\sigma(k)) - 1 + x_b - x_r > 0.$$

Therefore σ_{ik} is a swap such that $\Phi(\sigma_{ik}) < \Phi(\sigma)$.

We are left to prove that no segregated agents are created in σ_{ik} . The neighborhood of node $\sigma(k)$ in σ is composed of node $\sigma(i)$ and a remaining set of red nodes. Thus, when $\sigma(k)$ and $\sigma(i)$ exchange their colors in σ_{ik} , no segregated agents are created in the closed neighborhood of $\sigma(k)$. The neighborhood of node $\sigma(i)$ in σ is composed of node $\sigma(k)$ and a remaining set of x_b blue nodes and x_r red nodes (not counting $\sigma(k)$), with $x_b > x_r$. As the maximum degree of G is at most 3 and $\sigma(i)$ is adjacent to $\sigma(k)$, we have $x_b + x_r \leq 2$, which, since $x_b > x_r$, implies $x_r = 0$. Thus, when $\sigma(k)$ and $\sigma(i)$ exchange their colors in σ_{ik} , no segregated agents are created in the closed neighborhood of $\sigma(i)$. No other nodes are affected by the swap, thus no segregated agents are created. \square

Corollary 4. *There is a constant $c > 1$ such that, given a game (G, b, Λ) on a cubic graph G , with $\Lambda \leq 1/2$, the problem of computing an SE σ that maximizes $\text{Dol}(\sigma)$ is not c -approximable in polynomial time, unless $P = NP$.*

Proof. Let σ^* be a strategy profile that maximizes the value $\text{Dol}(\sigma^*)$. Thanks to Lemma 8, we know that there is an SE σ such that $\text{Dol}(\sigma) \geq \text{Dol}(\sigma^*)$. As a consequence, any SE that approximates $\text{Dol}(\sigma)$ within a factor of c would also approximate $\text{Dol}(\sigma^*)$ within a factor of c . The claim now follows from Theorem 15. \square

We now provide analogous results for bipartite graphs.

Theorem 16. *There is a constant $c > 1$ such that, given a game (G, b, Λ) on a bipartite graph G , the problem of computing a social optimum strategy profile is not c -approximable in polynomial time, unless $P = NP$.*

Proof. The reduction is from MINIMUM VERTEX COVER problem on cubic graphs, an optimization problem in which the goal is to compute a minimum-size set C of nodes of a given cubic graph G' such that every edge (u, v) of G' is covered by C , i.e., $\{u, v\} \cap C \neq \emptyset$. It is known that a minimum vertex cover on cubic graphs is not approximable within some constant $c' > 1$, unless $P = NP$, see [Alimonti and Kann, 1997].

Let us assume that n' and $m' = \frac{3}{2}n'$ is the number of nodes and edges of the input graph G' , respectively. We construct a graph G as follows (see Figure 7 for an example). G contains $n = n' + 7m' + 1$ nodes. More precisely, each node v of G' is modeled by a node x_v in G , while each edge e of G' is modeled by two nodes y_e^1 and y_e^2 in G . G also contains a special node z and $5m'$ additional dummy nodes. The special node z is connected by an edge to each of the $5m'$ dummy nodes and the n' nodes x_v , with v being a node of G' . Finally, G contains the two edges connecting x_v with y_e^1 and y_e^2 iff v is an endpoint of the edge e in G' .

By construction, we have that G is a bipartite graph. Let k^* denote the size of a minimum vertex cover of G' . We consider the game (G, b) played on the constructed graph G , where $b = k^* + 1$.

We claim that a social-optimal strategy profile σ^* has a $\text{Dol}(\sigma^*) = n$ iff G' admits a vertex cover of size k^* .

(\Leftarrow) Let C^* be a vertex cover of G' of size k^* . Consider the strategy profile σ^* in which one blue agent is placed on the special node z , while the remaining k^* blue agents are placed on the nodes x_v , with $v \in C^*$. Clearly, the red agents are placed on the remaining nodes of the graph. By construction, one can check that no agent in G is segregated (see Figure 7 for an example). Therefore, $\text{Dol}(\sigma^*) = n$.

(\Rightarrow) Let σ^* be a strategy profile such that $\text{Dol}(\sigma^*) = n$. First of all, as $k^* + 1 \leq n' + 1 < 5m'$, we have that no dummy node can be occupied by a blue agent. This is because all edges that connect a dummy node with the special node z must be bi-chromatic and the number of blue agents is not sufficient to cover all the dummy nodes. Therefore, all dummy nodes must be occupied by red agents and, as a consequence, the special node z is occupied by a blue agent. We claim that the set

$$C(\sigma^*) := \{v \in V(G') \mid x_v \text{ is occupied by a blue agent}\}$$

has size k^* and forms a vertex cover of G' . We observe that, by construction, it is enough to prove that $C(\sigma^*)$ has size k^* as each node y_e^i , with $i \in \{1, 2\}$, is adjacent to the nodes x_v

such that v covers e in G' . For the sake of contradiction, assume that $|C(\sigma^*)| < k$. We show the existence of a vertex cover of G' of size strictly smaller than k^* . Let E' be the subset of the edges of G' such that, for each $e \in E'$, y_e^1 and y_e^2 are both occupied by blue agents. Let C be a set of nodes of G' that contains $C(\sigma^*)$ plus any of the two endnodes of e , for each $e \in E'$. We now show that $|C| < k^*$. As $\text{Dol}(\sigma^*) = n$, each node y_e^i that is occupied by a red agent should be adjacent to a node x_v occupied by a blue agent. By construction v covers e and $v \in C(\sigma^*)$. Therefore, C covers all the edges of $E(G) \setminus E'$. Hence, C is a vertex cover of G' of size strictly smaller than k^* .

We complete the proof by showing that there is a constant $c > 1$ such that the problem of computing, in polynomial time, a strategy profile σ that approximates the social optimum σ^* is not approximable within c , unless $P = NP$. Let $c' > 1$ be the constant such that the MINIMUM VERTEX COVER problem on cubic graphs is not approximable within a factor of c' in polynomial time. As each node of G' covers 3 edges, we have that $k^* \geq \frac{m'}{3} = \frac{n'}{2}$. This implies that the MINIMUM VERTEX COVER problem on cubic graphs is approximable within a factor of 2 (all nodes of the graph suffice). Therefore, $c' < 2$. We set $c = 13/(14 - c')$. Observe that $c > 1$ as $1 < c' < 2$. We now prove that if we were able to compute, in polynomial time, a strategy profile σ such that $\text{Dol}(\sigma^*)/\text{Dol}(\sigma) \leq c$, then we could c' -approximate the MINIMUM VERTEX COVER problem on cubic graphs in polynomial time.

For the sake of contradiction, let σ be strategy profile that c -approximates $\text{Dol}(\sigma^*)$ and assume that σ can be computed in polynomial time. We use σ to define a new strategy profile σ' such that (i) $\text{Dol}(\sigma') \geq \text{Dol}(\sigma)$, (ii) one blue agent is placed on the special node z , and (iii) all the other blue agents are placed on a subset of nodes x_v with $v \in V(G')$.

First of all, we show that the special node z is occupied by a blue agent in σ . If not, there would be at least

$$5m' - k^* - 1 \geq 4m'$$

dummy nodes occupied by red agents and therefore $\text{Dol}(\sigma) \leq n - 4m'$. As a consequence, using also the fact that $m' = \frac{3}{2}n'$ and $m' \geq 1$, we would obtain

$$\begin{aligned} \frac{\text{Dol}(\sigma^*)}{\text{Dol}(\sigma)} &\geq \frac{n}{n - 4m'} = 1 + \frac{4m'}{n - 4m'} = 1 + \frac{4m'}{n' + 3m' + 1} \\ &\geq 1 + \frac{3m'}{\frac{2}{3}m' + 3m' + m'} = \frac{23}{14} > c, \end{aligned}$$

thus contradicting the fact that our solution is c -approximate.

The strategy profile σ' is obtained by modifying σ as follows. Blue agents that occupy dummy nodes exchange their position with red agents occupying nodes of the form x_v , with $v \in V(G')$. At the same time, every blue agent that occupies a node y_e^i , with $e \in E(G')$ and $i \in \{1, 2\}$, exchanges its position with a red agent occupying a node x_v such that $v \in V(G')$, where we give priority to nodes that cover e . Clearly, given σ , strategy profile σ' can be computed in polynomial time and we have $\text{Dol}(\sigma') \geq \text{Dol}(\sigma)$.

Let m'' be the number of edges of G' that are not covered by

$$C(\sigma') := \{v \in V(G') \mid x_v \text{ is occupied by a blue agent}\}.$$

We show that $m'' \leq (c' - 1)k^*$. First of all, we observe all the $5m'$ dummy nodes, the special node z and all the n' nodes x_v corresponding to the nodes $v \in V(G')$ are not segregated in σ' . Therefore, the number of uncovered edges of G' equals twice the number of segregated red agents in G (two red agents per uncovered edge e of G' that occupy the nodes y_e^1 and y_e^2). Therefore $\text{Dol}(\sigma') = n - 2m''$. As a consequence

$$c > \frac{\text{Dol}(\sigma^*)}{\text{Dol}(\sigma')} \geq \frac{n}{n - 2m''},$$

from which we derive

$$\begin{aligned} m'' &\leq \frac{c-1}{2c} n = \frac{\frac{13}{14-c'} - 1}{2 \frac{13}{14-c'}} n = \frac{c'-1}{26} (n' + 7m' + 1) \\ &\leq \frac{c'-1}{26} \left(\frac{2}{3}m' + 7m' + m' \right) = \frac{c'-1}{3} m' \\ &\leq (c'-1)k^*, \end{aligned}$$

where we use the facts that $m' = \frac{3}{2}n'$ (i.e., G' is cubic), $1 \leq m'$, and $k^* \geq \frac{m'}{3}$ (each node of G' covers 3 edges).

To complete the proof, let C be a vertex cover of G' that contains $C(\sigma')$ and a node that covers each of the edges of G' that are not covered by $C(\sigma')$. Clearly, given σ' , C can be computed in polynomial time. The size of C is upper bounded by the size of $C(\sigma')$ plus the number of uncovered edges, i.e.,

$$|C| \leq k^* + m'' \leq c'k^*.$$

Hence, C is a c' -approximate vertex cover of G' . This completes the proof. \square

Theorem 17. *There is a constant $c > 1$ such that, given a game $(G, b, 1/2)$ on a bipartite graph G , the problem of computing an SE σ that maximizes $\text{Dol}(\sigma)$ is not c -approximable in polynomial time, unless $P = NP$.*

Proof. We consider exactly the same reduction that we used in the proof of Theorem 16 and show the existence of a strategy profile σ^* that maximizes $\text{Dol}(\sigma^*)$ which is also an SE. Observe that, once we prove that σ^* is an SE, the rest of the proof can be derived from Theorem 16.

Consider the strategy profile σ^* in which a blue agent occupies the special node z , while the remaining k^* blue agents are placed on nodes of the form x_v such that $v \in C^*$ and C^* is an optimal vertex cover of G' (see also Figure 7). In the proof of Theorem 16 we already showed that $\text{Dol}(\sigma^*) = n$. In the following we show that σ^* is also a SE.

The red agents on the dummy nodes have maximum utility, so they never swap.

Let j be a red agent that is placed on a node of the form x_v , with $v \in V(G')$. The strategy σ_{ij} where i is a blue agent placed on a node x_v , with $v \in V(G')$, is not a profitable swap by Lemma 3. The strategy σ_{ij} where i is the blue agent placed on the special node z is not a profitable swap by Lemma 2.

Finally, consider any red agent j that is placed on a node of the form y_e^l with $e \in E(G')$ and $l \in \{1, 2\}$. This agent has a utility of either $p(1/3)$ or $p(2/3)$. But $p(1/3) = p(2/3)$

whenever $\Lambda \leq 1/2$. The strategy σ_{ij} , where i is a blue agent placed on a node x_v with $v \in V(G')$, is not a profitable swap either by Lemma 2 (when $\sigma(i)$ is adjacent to $\sigma(j)$) or simply because the utility of j in σ_{ij} is $p(7/8) < p(2/3)$. The blue agent i on the special node z has a utility that is strictly smaller than $p(1/6)$ as

$$\frac{k^* + 1}{5m' + n' + 1} \leq \frac{n'}{5n' + n' + 1} < 1/6.$$

Therefore, σ_{ij} is not a profitable swap because the utility of j in σ_{ij} is upper bounded by $p(5/6) < p(2/3)$. \square

7 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 [Agarwal *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., $\text{SUM}(\sigma) = \sum_{i=1}^n U_i(\sigma)$, 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

D. Acevedo-Garcia and K. A. Lochner. Residential segregation and health. *Neighborhoods and Health*, pages 265–87, 2003.

A. Agarwal, E. Elkind, J. Gan, A. Igarashi, W. Suksompong, and A. A. Voudouris. Schelling games on graphs. *Artif. Intell.*, 301:103576, 2021.

P. Alimonti and V. Kann. Hardness of approximating problems on cubic graphs. In *Algorithms and Complexity*, pages 288–298, 1997.

G. Barmpalias, R. Elwes, and A. Lewis-Pye. Digital morphogenesis via schelling segregation. In *FOCS 2014*, pages 156–165, 2014.

N. Betzler, A. Slinko, and J. Uhlmann. On the computation of fully proportional representation. *JAIR*, 47:475–519, 2013.

D. Bilò, V. Bilò, P. Lenzner, and L. Molitor. Topological influence and locality in swap schelling games. In *MFCS 2020*, pages 15:1–15:15, 2020.

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

N. Boehmer and E. Elkind. Individual-based stability in hedonic diversity games. In *AAAI 2020*, pages 1822–1829, 2020.

A. Bogomolnaia and M. O. Jackson. The stability of hedonic coalition structures. *GEB*, 38(2):201–230, 2002.

Ch. Brandt, N. Immorlica, G. Kamath, and R. Kleinberg. An analysis of one-dimensional schelling segregation. In *STOC 2012*, pages 789–804, 2012.

F. Brandt, M. Brill, E. Hemaspaandra, and L. A. Hemaspaandra. Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. *JAIR*, 53:439–496, 2015.

R. Bredereck, E. Elkind, and A. Igarashi. Hedonic diversity games. In *AAMAS 2019*, pages 565–573, 2019.

M. Bullinger, W. Suksompong, and A. A. Voudouris. Welfare guarantees in schelling segregation. *J. Artif. Intell. Res.*, 71:143–174, 2021.

D. Cable. The racial dot map. *Weldon Cooper Center for Public Service, University of Virginia*, 2013.

H. Chan, M. T. Irfan, and C. V. Than. Schelling models with localized social influence: A game-theoretic framework. In *AAMAS 2020*, pages 240–248, 2020.

A. Chauhan, P. Lenzner, and L. Molitor. Schelling segregation with strategic agents. In *SAGT 2018*, pages 137–149, 2018.

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

H. Echezell, T. Friedrich, P. Lenzner, L. Molitor, M. Pappik, F. Schöne, F. Sommer, and D. Stangl. Convergence and hardness of strategic schelling segregation. In *WINE 2019*, pages 156–170, 2019.

E. Elkind, P. Faliszewski, and P. Skowron. A characterization of the single-peaked single-crossing domain. In *AAAI 2014*, 2014.

N. Immorlica, R. Kleinberg, B. Lucier, and M. Zadomighad-dam. Exponential segregation in a two-dimensional schelling model with tolerant individuals. In *SODA 2017*, pages 984–993, 2017.

S. A. Jackson, R. T. Anderson, N. J. Johnson, and P. 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.

- P. Kanellopoulos, M. Kyropoulou, and A. A. Voudouris. Modified schelling games. *TCS*, 880:1–19, 2021.
- P. Kanellopoulos, M. Kyropoulou, and A. A. Voudouris. Not all strangers are the same: The impact of tolerance in schelling games. *CoRR*, abs/2105.02699, 2021.
- L. Kreisel, N. Boehmer, V. Froese, and R. Niedermeier. Equilibria in schelling games: Computational complexity and robustness. *CoRR*, abs/2105.06561, 2021.
- D. S. Massey and N. A. Denton. The dimensions of residential segregation. *Soc. Forces*, 67(2):281–315, 1988.
- D. S. Massey and N. A. Denton. *American Apartheid: Segregation and the Making of the Underclass*. Harvard U. Press, 1993.
- D. Monderer and L. S. Shapley. Potential games. *GEB*, 14(1):124–143, 1996.
- Th. C. Schelling. Dynamic models of segregation. *J. Math. Soc.*, 1(2):143–186, 1971.
- Th. C. Schelling. *Micromotives and Macrobehavior*. WW Norton & Company, 2006.
- Tom W. Smith, M. Davern, J. Freese, and S. L. Morgan. General social surveys, 1972–2018 cumulative codebook. *NORC ed. Chicago: NORC 2019, U. Chicago*, 2019.
- V. V. Vazirani. *Approximation algorithms*. Springer Science & Business Media, 2013.
- T. Walsh. Uncertainty in preference elicitation and aggregation. In *AAAI 2007*, pages 3–8, 2007.
- L. Yu, H. Chan, and E. Elkind. Multiwinner elections under preferences that are single-peaked on a tree. In *AAAI 2013*, pages 425–431, 2013.
- J. Zhang. A dynamic model of residential segregation. *J. Mat. Soc.*, 28(3):147–170, 2004.
- J. Zhang. Residential segregation in an all-integrationist world. *J. Econ. Behav. and Org.*, 54(4):533–550, 2004.