

The Bakers and Millers Game with Restricted Locations

Simon Krogmann
Hasso Plattner Institute,
University of Potsdam
Potsdam, Germany
simon.krogmann@hpi.de

Pascal Lenzner
Institute of Computer Science,
University of Augsburg
Augsburg, Germany
pascal.lenzner@uni-a.de

Alexander Skopalik
Mathematics of Operations Research,
University of Twente
Enschede, The Netherlands
a.skopalik@utwente.nl

ABSTRACT

We study strategic location choice by customers and sellers, termed the Bakers and Millers Game in the literature. In our generalized setting, each miller can freely choose any location for setting up a mill, while each baker is restricted in the choice of location for setting up a bakery. For optimal bargaining power, a baker would like to select a location with many millers to buy flour from and with little competition from other bakers. Likewise, a miller aims for a location with many bakers and few competing millers. Thus, both types of agents choose locations to optimize the ratio of agents of opposite type divided by agents of the same type at their chosen location. Originally raised in the context of Fractional Hedonic Games, the Bakers and Millers Game has applications that range from commerce to product design.

We study the impact of location restrictions on the properties of the game. While pure Nash equilibria trivially exist in the setting without location restrictions, we show via a sophisticated, efficient algorithm that even the more challenging restricted setting admits equilibria. Moreover, the computed equilibrium approximates the optimal social welfare by a factor of at most $2 \left(\frac{e}{e-1} \right)$. Furthermore, we give tight bounds on the price of anarchy/stability.

On the conceptual side, the location choice feature adds a new layer to Hedonic Games, in the sense that agents that select the same location form a coalition. This allows to naturally restrict the possible coalitions that can be formed. With this, our model generalizes simple symmetric Fractional Hedonic Games on complete bipartite valuation graphs and also Hedonic Diversity Games with utilities single-peaked at 0. We believe that this generalization is also a very interesting direction for other types of Hedonic Games.

KEYWORDS

Competitive Location Choice; Fractional Hedonic Games; Nash Equilibrium; Price of Anarchy

ACM Reference Format:

Simon Krogmann, Pascal Lenzner, and Alexander Skopalik. 2025. The Bakers and Millers Game with Restricted Locations. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025)*, Detroit, Michigan, USA, May 19 – 23, 2025, IFAAMAS, 9 pages.

1 INTRODUCTION

Markets facilitate trade. They are the key places where supply meets demand. Typically, there are many possible markets for goods or

services, and thus sellers face the economic decision of which markets to serve, while customers strategically decide which market to patronize. The choice of the right location for trade is crucial in a competitive economic environment both for sellers and for customers. The number of competitors in a given market is a key factor in determining supply and demand, and eventually, prices and profits from trade. For achieving high prices, sellers want to minimize competition with other sellers and at the same time, they want to maximize the number of customers. For customers who aim for low prices, it is important to patronize a market with many sellers to choose from and with low competition from other customers.

Our running example of the above competitive setting will be the Bakers and Millers Game, as briefly introduced by Aziz, Brandt, and Harrenstein [4]. There, millers choose locations to set up their mill to produce flour whereas bakers depend on buying their flour from local millers. In this scenario, millers aim for locations with many bakers as customers and only a few other competing millers. Bakers will patronize a local miller that has as few competing customers as possible. If both bakers and millers can freely choose any location, then this setting is well-studied, since it corresponds to a simple symmetric Fractional Hedonic Game on a complete bipartite valuation graph [3, 4]. However, in many practical applications, there is an asymmetry between sellers and customers, i.e., sellers can freely choose where to set up their stores, but customers can only patronize a restricted set of locations, e.g., their local neighborhood. In this paper, we study the impact of such an asymmetry of location choice by investigating the Bakers and Millers Game where one side has restricted locations. In our setting, the millers can freely choose their location, while the bakers can only choose from an agent-specific restricted set of locations. See Figure 1 for an illustration. With this, we investigate a very simple model that




Figure 1: Illustration of the Bakers and Millers Game with restricted locations. Bakers choose between an agent-specific set of locations (black arrows mark the selected location, dashed lines the unselected ones), while millers can choose any location. On the left, miller m can improve her utility by choosing location y for utility $\frac{3}{2}$ or location z for utility 1. A pure Nash equilibrium is given on the right.



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

naturally extends simple symmetric Fractional Hedonic Games on complete bipartite valuation graphs.

More examples are classical farmer's markets but also online commerce on different platforms like Alibaba, Amazon, eBay or Zalando. Another application is gaining a favorable position in a two-sided matching market: A worker wants to move to a city that has more companies and thus more possible jobs to choose from, whereas a company wants to move to a city which has fewer companies that compete for the local workforce. However, location choice is not restricted to trade. "Sellers" could be product designers, politicians, researchers, or writers who strategically focus on certain topics, and the "customers" might be consumers, voters, reviewers or readers only interested in some of these topics. Thus, "locations" might simply be points in some abstract feature space.

Given the importance of strategic location choice, the research on such problems dates back roughly a century, starting with Hotelling's seminal model of several sellers competitively selecting a location for their store on the city's main street [30] that was later adopted by Downs for modeling political competition [15]. Based on this, many other location choice models have been investigated in Economics, Mathematics, Computer Science and Operations Research and this field is typically called Location Analysis [19, 43]. However, most of these models solely focus on the seller's side, i.e., on locating stores in settings where the customers do not face strategic decisions since they simply patronize their nearest store. Also, the opposite point of view has been taken, where the sellers have fixed positions and the focus is on the customer behavior, like in consumer behavior theory [8, 45].

The setting of rational agents who compete over a fixed set of offered services, products or resources is covered by the broad area of Congestion Games [39, 44]. There, agents strategically select resources and an agent's utility for a given resource deteriorates with the number of other agents who also choose the same resource. Thus, in the Bakers and Millers Game, both bakers and millers face a Singleton Congestion Game if the behavior of the other type of agents is fixed, i.e., given the location choice of the millers, the bakers want to select locations with a minimum number of other bakers and a maximum number of millers and vice versa. Hence, our game combines both the sellers' and the customers' perspectives simultaneously, i.e., both sides strategically select a location which then determines which customers are matched to which sellers. Our game focuses on the impact of the location choice and it abstracts away from the exact mechanisms that determine prices via supply and demand. This allows for modeling a broad range of applications ranging from commerce, over two-sided matching, to the design of products or political campaigns.

1.1 Our Contribution

We study competitive location choice in the Bakers and Millers Game introduced by Aziz, Brandt, and Harrenstein [4]. Most importantly, we generalize this setting by considering that one side, the bakers, may have location restrictions, i.e., they can only choose their location from an individually predetermined subset of possible locations. This naturally models an asymmetry in many competitive settings where sellers or producers can freely choose where to open stores or facilities while the customers are bound to patronize

stores accessible to them only, e.g., local stores in their vicinity. Besides store locations, this can also model product or feature choices by producers in a setting where the customers are only interested in certain subsets of features or products, like family-friendly cars, lightweight e-bikes, or regional vegan food.

The Bakers and Millers Game with restricted locations generalizes simple symmetric Fractional Hedonic Games on complete bipartite valuation graphs by naturally limiting the possible coalitions that can be formed. The same holds for Hedonic Diversity Games with utilities that have a single peak at 0. As these two examples indicate, we emphasize that this extra layer of first selecting locations that then determine the formed coalitions can be added to any variant of Hedonic Games to better model real-world settings.

As our main contribution, we analyze the impact of the location restrictions in the Bakers and Millers Game. While proving the existence of pure Nash equilibria is trivial if both bakers and millers can freely choose their locations, this task is much more challenging if the bakers have restricted locations. Despite this, we show the existence of pure Nash equilibria by giving a polynomial-time algorithm to compute one. As a consequence of the location restrictions, some bakers might end up without an available miller. Hence, we use the *coverage of bakers*, i.e., the number of bakers with at least one available miller, as our social welfare function. For this function, our aforementioned efficient algorithm computes a pure Nash equilibrium that approximates the optimal social welfare by $\left(1 + \frac{\min(|\mathcal{L}|, |\mathcal{M}|-1)}{|\mathcal{M}|}\right) \frac{e}{e-1} < 2 \left(\frac{e}{e-1}\right)$ with Euler's number e and the sets of locations $|\mathcal{L}|$ and millers $|\mathcal{M}|$. On the other hand, we show that computing a strategy profile or even a pure Nash equilibrium with optimal social welfare is NP-hard. Finally, we establish tight bounds of $1 + \frac{\min(|\mathcal{L}|, |\mathcal{M}|-1)}{|\mathcal{M}|} < 2$ on the price of stability and of $|\mathcal{B}|$ on the price of anarchy, with the set of bakers \mathcal{B} .

1.2 Preliminaries and Model

Basics. We denote the set of available locations by \mathcal{L} and the set of agents by $\mathcal{A} = \mathcal{B} \cup \mathcal{M}$, with the set of bakers \mathcal{B} and millers \mathcal{M} . Each baker $b \in \mathcal{B}$ has a set of feasible locations $L(b) \subseteq \mathcal{L}$.

A strategy profile is given by the tuple (\mathbf{s}, \mathbf{t}) , where \mathbf{s} is a vector of bakers' locations $(s_b)_{b \in \mathcal{B}}$, with $s_b \in L(b)$, and \mathbf{t} is a vector of millers' locations $(t_m)_{m \in \mathcal{M}}$, with $t_m \in \mathcal{L}$. We denote by $B_s(\ell) := |\{b \in \mathcal{B} \mid s_b = \ell\}|$ the number of bakers at location ℓ for a given strategy profile \mathbf{s} . Similarly, $M_t(\ell) := |\{m \in \mathcal{M} \mid t_m = \ell\}|$ denotes the number of millers at location ℓ .

Objectives. Agents choose locations to avoid competition due to agents of the same type while maximizing the availability of agents of the other type. In its simplest form, this notion is captured by a utility function that corresponds to the ratio of the two types of agents. Hence, the (always well-defined) utility of a baker b and a miller m in a given strategy profile (\mathbf{s}, \mathbf{t}) is

$$u_b(\mathbf{s}, \mathbf{t}) = \frac{M_t(s_b)}{B_s(s_b)} \quad \text{and} \quad u_m(\mathbf{s}, \mathbf{t}) = \frac{B_s(t_m)}{M_t(t_m)}, \text{ respectively.}$$

Equilibria. A state $(\mathbf{s}^*, \mathbf{t}^*)$ is a *miller equilibrium* if for each miller $m \in \mathcal{M}$ there is no location $t'_m \in \mathcal{L}$ such that

$$u_m(\mathbf{s}^*, (t'_m, \mathbf{t}_{-m}^*)) > u_m(\mathbf{s}^*, \mathbf{t}^*).$$

A state (s^*, t^*) is a *baker equilibrium* if for each baker $b \in \mathcal{B}$ there is no location $s'_b \in L(b)$ such that

$$u_b((s'_b, s^*_{-b}), t^*) > u_b(s^*, t^*).$$

A state (s^*, t^*) is a pure Nash equilibrium if it is both a baker equilibrium and a miller equilibrium.

Efficiency. We measure the efficiency of equilibria by the number of bakers that can shop at some miller. Therefore, we define the *coverage* of a strategy profile (s, t) as the number of bakers with at least one miller at the same location, i.e.,

$$W(s, t) = |\{b \in \mathcal{B} \mid M_t(s_b) > 0\}|.$$

In Section 3, we justify this definition. Based on the coverage, we define the *price of anarchy* [32] and the *price of stability* [2] as

$$\text{PoA} = \sup_I \left(\frac{W(\text{OPT}(I))}{W(\text{worstNE}(I))} \right) \quad \text{and} \quad \text{PoS} = \sup_I \left(\frac{W(\text{OPT}(I))}{W(\text{bestNE}(I))} \right),$$

where for an instance I , we have $\text{OPT}(I)$ as the social optimum, $\text{worstNE}(I)$ as the Nash equilibrium with the lowest coverage and $\text{bestNE}(I)$ as the Nash equilibrium with the highest coverage.

1.3 Related Work

Research in the broad field of Hedonic Games was initiated by the seminal works of Dreze and Greenberg [16] and Bogomolnaia and Jackson [10]. In these games, the agents have preferences over all possible coalitions of agents and every partition of the set of agents into disjoint coalitions is considered an outcome. The crucial aspect of Hedonic Games is that the utility of an agent in some coalition only depends on the composition of her coalition. Given this, many variants of utility functions have been studied.

Fractional Hedonic Games. Closest to our model are Fractional Hedonic Games (FHGs) introduced by Aziz, Brandt, and Harrenstein [4] and later extended [3], where each agent has an individual value for every other agent and the utility of some agent in a coalition is her average agent value within her coalition, i.e., it is the sum over the individual values of all the agents in her coalition divided by the number of agents in the coalition.

A nice feature of FHGs is that an instance can be described via a weighted directed *valuation graph*, where an edge (u, v) with weight w encodes that agent u has a value of w for agent v . A FHG is called *symmetric* if an edge (u, v) with weight w in the corresponding valuation graph implies the existence of the reverse edge (v, u) with the same weight w , i.e., pairs of agents value each other equally. A FHG is called *simple* if all edge weights in the valuation graph are either 0 or 1. Thus, the valuations in simple symmetric FHGs can be modeled with an undirected unweighted valuation graph.

If the number of locations in our game equals the total number of bakers and millers and if the bakers can choose any location, then our game is equivalent to a simple symmetric FHG. On the other hand, if the valuation graph of a simple symmetric FHG is bipartite and complete, then this is equivalent to our game¹. Thus, our model generalizes simple symmetric FHGs that have bipartite and complete valuation graphs.

¹FHGs and our game differ slightly in the definition of the agents' valuations. However, it is easy to show that the agents' preferences over coalitions/locations are equivalent.

FHGs have been intensively studied with different notions of stability, like core, Nash or individual stability [3, 6, 10, 11, 13, 36] and also in a modified version where the agents do not count themselves in the computation of their utility in a coalition [37, 38]. In our paper, we use Nash stability. In a Nash stable partition, no agent can improve her utility by unilaterally changing her coalition.

Nash stable partitions are not guaranteed to exist for FHGs with arbitrary agent valuations, but they do exist for non-negative valuations [6], which include simple FHGs. Also, it is NP-complete to decide if a given instance admits a Nash stable partition, even with symmetric valuations [11]. If the quality of stable partitions is measured with the utilitarian welfare, then the price of anarchy is in $\Theta(n)$, even on unweighted paths, and the price of stability on weighted graphs is in $\Theta(n)$, even on weighted stars [6]. Note that we also prove bounds on the price of anarchy and stability but with a different social welfare function. Bilò, Fanelli, Flammini, Monaco, and Moscardelli [6] show that best response dynamics are not guaranteed to converge to a Nash stable partition even on unweighted bipartite graphs and that computing a Nash stable partition with maximum social welfare is NP-hard. However, the non-convergence result does not carry over to our game since the bipartite valuation graph in [6] is not complete.

Hedonic Diversity Games. Also related are Hedonic Diversity Games (HDGs) [9, 12]. There, the agents have different types and their utility depends on the type-composition of their coalition. Similarly to FHGs, the utility of an agent is the fraction of own-type agents. The utilities in our game correspond to single-peaked utilities, a concept that dates back to Black [7], and such preferences have also been investigated in HDGs [9, 12].

As with FHGs, our game is equivalent to HDGs with utilities that are single-peaked at 0 if the number of locations equals the total number of bakers and millers and if all bakers can choose any location. Thus, by restricting the number of locations or the access to these locations, our game generalizes such HDGs.

Schelling Games and Resource Selection. Another class of related games are Schelling Games [1, 14, 18]. There, agents of different types select a location on a given graph and the utility of an agent is a function of the type composition of the neighborhood on her selected location. The crucial difference to FHGs and HDGs is that these neighborhoods may be non-overlapping. The agents' utility is a threshold function that attains its maximum value if the fraction of same-type agents in the neighborhood is at least τ , for some $\tau \in [0, 1]$. Recently, single-peaked utility functions have also been studied [5, 25].

Close to Schelling Games and to our game is the model presented by Gadea Harder et al. [26]. In their setting, agents of different types jointly select resources. Like in Schelling Games, an agent is content with her resource if the fraction of same-type agents selecting the resource reaches at least a tolerance threshold $\tau \in [0, 1]$. Two variants are studied: impact-blind and impact-aware agents. For the former, equilibrium existence is shown via a potential function argument while for the latter, this only holds for $\tau \leq \frac{1}{2}$.

By suitably modifying the utilities and assuming $\tau = 1$, it can be shown that our Bakers and Millers Game is a special case of the model with impact-aware agents in [26] with inverted utilities. However, since their potential argument only works for $\tau \leq \frac{1}{2}$, their

existence result does not carry over to our game. Also, in the other direction, our positive results do not carry over to their model.

Strategic Facility Location. Finally, strategic facility location models [41] are also related, in particular, the model by Vetta [47], where sellers strategically select markets with a certain purchasing power and the price is explicitly determined by the delivery costs. Voronoi Games [17, 22], Location Games on Networks [24], Market Sharing Games [27], and Network Investment Games [46] are also similar since agents strategically select a location or resources to maximize their utility. Closer to our setting are models where the clients also face a strategic choice in the form of minimizing a weighted sum of their travel and waiting time [23, 31, 42]. Recently, this setting with both strategic facilities and strategic clients was studied as a sequential game [33–35], where first sellers select a location on a given graph and then the customers on graph nodes decide how to distribute their purchasing power among neighboring sellers. Our model resembles a simultaneous variant of this game with different utility functions.

2 EQUILIBRIUM EXISTENCE AND COMPUTATION

We observe that equilibria in the Bakers and Millers Game are not unique and can be quite different. See Figure 2 for an example with two locations x and y , four bakers and two millers. In the left equilibrium, all agents except one baker are at the same location x . On the right, each location has one miller and two bakers.




Figure 2: An instance with two Nash equilibria which differ in the agents' utilities and the coverage. Of the bakers, only b has both locations available in $L(b)$.

In the special case of $L(b) = \mathcal{L}$, i.e., where each baker can choose every location, there is always a trivial pure Nash equilibrium by assigning all agents to a single location. Thus, we now turn to the non-trivial case with restricted bakers and devise an algorithm to compute a pure Nash equilibrium. This algorithm, shown in Algorithm 1, proceeds in three phases: First, it tries to concentrate the bakers as much as possible in few locations so that they are placed in the fullest location they can access. Then, a miller equilibrium is created by sequentially inserting each miller at the currently best location. Finally, the bakers are redistributed to a specific baker equilibrium, which maintains the miller equilibrium. We show correctness in the following proof with two embedded lemmas.

THEOREM 2.1 (NASH EQUILIBRIUM ALGORITHM). *Algorithm 1 computes a pure Nash equilibrium in polynomial time.*

PROOF. Algorithm 1 first greedily determines the order of the locations by the number of yet unassigned bakers in range and

Algorithm 1: compute an equilibrium profile (s^*, t^*)

```

1  $\mathcal{L}' \leftarrow \mathcal{L}, \mathcal{B}' \leftarrow \mathcal{B};$ 
2 for  $i = 1$  to  $|\mathcal{L}'|$  do
3    $\ell_i \leftarrow \arg \max_{\ell \in \mathcal{L}'} |\{b \in \mathcal{B}' \mid \ell \in L(b)\}|$ , a location with
     most bakers in range;
4   for baker  $b \in \mathcal{B}'$  with  $\ell_i \in L(b)$  do
5      $s_b \leftarrow \ell_i$ ;
6    $\mathcal{L}' \leftarrow \mathcal{L}' \setminus \{\ell_i\}, \mathcal{B}' \leftarrow \mathcal{B}' \setminus \{b \mid \ell_i \in L(b)\};$ 
7 for miller  $m \in \mathcal{M}$  do
8    $t_m^* \leftarrow$  location  $\ell$  (with smallest index) maximizing  $\frac{B_s(\ell)}{M_i(\ell)+1}$ ;
9  $s^* \leftarrow$  baker strategy profile maximizing  $\Phi_{t^*}(s^*)$  (Lemma 2.4);
```




Figure 3: Illustration of the first loop of Algorithm 1: Location x has the most bakers in its range, so it is chosen first as ℓ_1 and assigned all possible bakers. Of the remaining bakers, location z has the most in its range. Finally, there are no bakers left for location y . Observe that the bakers can never move to a location that was removed earlier (i.e. that is to the left in the ordered result in (b)).

assigns the bakers to those locations in that order resulting in a baker profile s . See Figure 3 for an example. We then iteratively insert millers to their best-response locations. This yields a profile (s, t^*) in which the millers are in equilibrium. Finally, we compute a baker equilibrium profile s^* by choosing the state s that maximizes the classic Rosenthal potential function [44]

$$\Phi_{t^*}(s) = \sum_{\ell \in \mathcal{L}} M_{t^*}(\ell) \cdot H_{B_s(\ell)},$$

where H_i denotes the i -th harmonic number. We will show that the resulting profile (s^*, t^*) remains an equilibrium for the millers, too. We begin by proving that after inserting the millers in Line 8, the millers are indeed in a miller equilibrium.

LEMMA 2.2 (MILLER EQUILIBRIUM). *In (s, t^*) every miller is allocated to her best-response location.*

PROOF. Assume for the sake of contradiction there is a miller m on location ℓ that could improve to a location ℓ' , i.e.,

$$\frac{B_s(\ell)}{M_{t^*}(\ell)} < \frac{B_s(\ell')}{M_{t^*}(\ell') + 1}. \quad (1)$$

Consider the last miller that was placed on location ℓ . For the number x of millers located at ℓ' at this time, we have

$$\frac{B_s(\ell)}{M_{t^*}(\ell)} \geq \frac{B_s(\ell')}{x + 1} \quad (2)$$

at this time. Otherwise, m would have been placed on ℓ' instead. As $x \leq M_{t^*}(\ell')$, the lemma follows. \square

Before we show that the last step does maintain the property that millers are in equilibrium, let us remark some useful observations. The number of assigned bakers in s is non-increasing with the index of the location, i.e., $B_s(\ell_i) \geq B_s(\ell_{i+1})$, and each baker b is assigned to the location $\ell_i \in L(b)$ with smallest i of locations in $L(b)$. This can be seen in Figure 3b. Therefore, when moving from s to s^* , bakers are only moved to locations with higher indices. Furthermore, by inserting the millers at their best-response locations and breaking ties in favor of smaller indices, we have that $M_{t^*}(\ell_i) \geq M_{t^*}(\ell_{i+1})$.

LEMMA 2.3 (MILLER EQUILIBRIUM MAINTAINED). *Let t^* be a miller profile computed by Algorithm 1 and $s^* = \arg \max_x \Phi_{t^*}(x)$. The state (s^*, t^*) is a miller equilibrium.*

PROOF. To prove the lemma, we analyze the changes when moving from s to s^* . In particular, we show that the minimal utility of any miller does not decrease and that the maximal achievable utility of a deviation does not increase. To that end, we denote by $L_{t^*} \subseteq \mathcal{L}$ the subset of locations occupied by millers in t^* . We consider the *difference graph* which is a multi-graph $G = (L_{t^*}, E)$ that describes the change of bakers from s to s^* . See Figure 4 for an example. Therefore, E contains an arc from location ℓ_i to location ℓ_j for each baker b with $s_b = \ell_i$ and $s_b^* = \ell_j$. Note that we always have $i < j$ and, therefore, the difference graph G is acyclic.




Figure 4: For Lemma 2.3 and the state (s, t^*) on the left, the arrows indicate the strategy changes of bakers from s to s^* , e.g., there is a baker b with $s_b = v$ and $s_b^* = z$. This results in the difference graph G on the right, with the two paths P_1 and P_2 making up \mathcal{P} .

We consider a path decomposition \mathcal{P} of G into maximal paths. In particular, there is no location that is both a start and an endpoint of paths, as these two paths would be merged. Let $(p_1, p_2, \dots, p_m) \in \mathcal{P}$ be a path consisting of locations. As s^* maximizes the function $\Phi_{t^*}(x) = \sum_{\ell \in \mathcal{L}} M_{t^*}(\ell) \cdot H_{B_x}(\ell)$, we have that

$$\frac{M_{t^*}(p_1)}{B_{s^*}(p_1) + 1} \leq \frac{M_{t^*}(p_m)}{B_{s^*}(p_m)} \quad \text{or equiv.} \quad \frac{B_{s^*}(p_1) + 1}{M_{t^*}(p_1)} \geq \frac{B_{s^*}(p_m)}{M_{t^*}(p_m)} \quad (3)$$

since otherwise rearranging the bakers backwards along that path would increase the value of Φ . Note that rearranging changes only the number of bakers on p_1 and p_m .

By the choice of paths in \mathcal{P} being *maximal*, we have that p_1 has only outgoing arcs and p_m has only incoming arcs in G . Therefore,

$$B_{s^*}(p_1) + 1 \leq B_s(p_1) \quad \text{and} \quad B_{s^*}(p_m) \geq B_s(p_m) + 1. \quad (4)$$

Combining Equations (3) and (4) we get that the maximal utility of any miller in s^* is not larger than in s :

$$\frac{B_s(p_1)}{M_{t^*}(p_1)} \geq \frac{B_{s^*}(p_1) + 1}{M_{t^*}(p_1)} \geq \frac{B_{s^*}(p_m)}{M_{t^*}(p_m)} \geq \frac{B_s(p_m) + 1}{M_{t^*}(p_m)}. \quad (5)$$

Since we have $M_{t^*}(p_1) \geq M_{t^*}(p_m)$ because p_1 has a lower index in $(l_1, \dots, l_{|\mathcal{L}|})$ than p_m , we also have

$$\frac{B_{s^*}(p_1)}{M_{t^*}(p_1)} \geq \frac{B_s(p_m)}{M_{t^*}(p_m)} \quad (6)$$

which implies that the minimal utility of any miller in s^* is not smaller than in s . Now assume for the sake of contradiction that there is a miller m that has a profitable deviation from location $\ell = t^*(m)$ to location ℓ' in (s^*, t^*) . Furthermore, let m' be the miller with minimal utility u_{\min} in (s^*, t) . Therefore, we have that

$$\frac{B_{s^*}(\ell')}{M_{t^*}(\ell') + 1} \geq \frac{B_{s^*}(\ell)}{M_{t^*}(\ell)} \geq u_{\min} \quad (7)$$

where the second inequality follows from Equation (6). However, by Lemma 2.2 this deviation for m' was not profitable in (s, t^*) :

$$\frac{B_s(\ell')}{M_{t^*}(\ell') + 1} \leq u_{\min}. \quad (8)$$

Combining Equations (7) and (8) implies that $B_{s^*}(\ell') > B_s(\ell')$, so ℓ' is the endpoint of some path in \mathcal{P} . We now consider a path in the difference graph G that ends at ℓ' and starts from some location ℓ_{start} . By Equation (5) we have that

$$\frac{B_s(\ell_{\text{start}})}{M_{t^*}(\ell_{\text{start}})} \geq \frac{B_{s^*}(\ell')}{M_{t^*}(\ell')}. \quad (9)$$

By the existence of a path from ℓ' to ℓ_{start} , we get $M_{t^*}(\ell_{\text{start}}) \geq M_{t^*}(\ell')$. Therefore, we have

$$\frac{B_s(\ell_{\text{start}})}{M_{t^*}(\ell_{\text{start}}) + 1} \geq \frac{B_{s^*}(\ell')}{M_{t^*}(\ell') + 1}. \quad (10)$$

However, then deviating to ℓ_{start} would already have been an improving deviation for m' in (s, t^*) , which contradicts Lemma 2.2. Therefore, if there is no profitable unilateral deviation for a miller in s , there is also none in s^* , which concludes the lemma. \square

Finally, it remains to show that (s^*, t^*) is a baker equilibrium, which follows from the fact that Φ is an exact potential function for the induced baker game if we fix all millers. Assume for the sake of contradiction that there is a baker b that has a profitable deviation from $\ell = s_b^*$ to ℓ' , with

$$\frac{M_{t^*}(\ell)}{B_{s^*}(\ell)} < \frac{M_{t^*}(\ell')}{B_{s^*}(\ell') + 1}. \quad (11)$$

But then

$$\Phi_{t^*}((s_{-b}^*, \ell')) = \Phi_{t^*}(s^*) - \frac{M_{t^*}(\ell)}{B_{s^*}(\ell)} + \frac{M_{t^*}(\ell')}{B_{s^*}(\ell') + 1} > \Phi_{t^*}(s^*),$$

which contradicts the choice of s^* as a maximizer of Φ_{t^*} . In the following Lemma 2.4, we show how to compute the maximizer in polynomial time and thus our current Theorem 2.1 follows. \square

To compute the maximizer in Line 9 of Algorithm 1, we employ a reduction to integral MinCostFlow similar to the algorithm of Fabrikant, Papadimitriou, and Talwar [20, Theorem 2] for symmetric congestion games.




Figure 5: To compute the maximizer of Φ_{t^*} , the instance on the left with the given assignment of millers t^* is reduced to the MINCOSTFLOW instance on the right by Lemma 2.4. All edge capacities are 1 and the costs are given by the labels.

LEMMA 2.4 (POTENTIAL FUNCTION MAXIMIZER). *The maximizer of $\Phi_{t^*} = \sum_{\ell \in \mathcal{L}} M_{t^*}(\ell) \cdot H_{B_s(\ell)}$ can be computed in polynomial time.*

PROOF. The set of nodes of the MINCOSTFLOW instance corresponds to the union of the set of bakers \mathcal{B} , the set of locations \mathcal{L} , and the set of a source s and a sink t . There are edges with capacity 1 and cost 0 from s to all bakers and from each baker to her accessible locations. From each location ℓ , there are $|\mathcal{B}|$ parallel edges to t with capacity 1 and costs $-\frac{M_{t^*}(\ell)}{1}, -\frac{M_{t^*}(\ell)}{2}, \dots, -\frac{M_{t^*}(\ell)}{|\mathcal{B}|}$. We give an example of this construction in Figure 5.

Since in any integral MINCOSTFLOW for a given location only the edges with minimal cost are used, the cost of a flow is equal to $-\Phi_{t^*}(s) = -\sum_{\ell \in \mathcal{L}} M_{t^*}(\ell) \cdot H_{B_s(\ell)}$ for the corresponding baker assignment s . Thus, an integral MINCOSTFLOW in this network corresponds to an assignment of bakers to locations that maximizes the function Φ_{t^*} .

Note that our instance contains negative costs but no negative cycles. Therefore, we may use Orlin’s algorithm [40]. \square

We observe that the runtime of Algorithm 1 is dominated by the computation of the MINCOSTFLOW. The construction yields an instance with $O(|\mathcal{B}| + |\mathcal{L}|)$ nodes and $O(|\mathcal{B}||\mathcal{L}|)$ edges. In Section 3.2, we will give a bound on the social welfare approximation of the algorithm.

3 SOCIAL WELFARE

The most prominent notion of social welfare is utilitarian social welfare, i.e., the summation of the agents’ utilities. However, in the game at hand, this sum essentially simplifies to the number of millers plus the number of covered bakers, where the latter denotes bakers for which there is a miller on a location they can choose:

$$\begin{aligned} \text{Welfare} &= \sum_{b \in \mathcal{B}} u_b(s, t) + \sum_{m \in \mathcal{M}} u_m(s, t) \\ &= \sum_{\ell: B_s(\ell) \neq \emptyset} B_s(\ell) \frac{M_t(\ell)}{B_s(\ell)} + \sum_{\ell: M_t(\ell) \neq \emptyset} M_t(\ell) \frac{B_s(\ell)}{M_t(\ell)} \\ &= \sum_{\ell: B_s(\ell) \neq \emptyset} M_t(\ell) + \sum_{\ell: M_t(\ell) \neq \emptyset} B_s(\ell). \end{aligned} \quad (12)$$

Note that in equilibrium, the first term of Equation (12) evaluates to $|\mathcal{M}|$ and the second term is equal to the number of covered bakers. We therefore define the latter as *coverage* and use it as a metric for social welfare in the rest of the paper.

Definition 3.1 (Coverage). The coverage of a strategy profile (s, t) is the number of bakers with at least one miller at the same location:

$$W(s, t) = |\{b \in \mathcal{B} \mid M_t(s_b) > 0\}|.$$

For this measure, we show that the price of anarchy is linear in $|\mathcal{B}|$, while the price of stability is $1 + \frac{\min(|\mathcal{L}|, |\mathcal{M}|-1)}{|\mathcal{M}|}$. We also show that computing both the social welfare optimum and the optimal Nash equilibrium is NP-hard. On the positive side, we show that our Algorithm 1 actually computes a Nash equilibrium that approximates the optimal social welfare by a factor of $(1 + \frac{\min(|\mathcal{L}|, |\mathcal{M}|-1)}{|\mathcal{M}|}) \frac{e}{e-1}$.

3.1 Price of Anarchy and Stability

We show that the price of anarchy is $|\mathcal{B}|$, which is the worst possible in terms of the number of bakers. This follows from the fact that in every instance there exists a Nash equilibrium which covers at least one baker, while for some instances there is an equilibrium which covers exactly one baker.




Figure 6: This instance illustrates the worst case for the price of anarchy. The bakers are restricted to x and the location directly below them while the miller is unrestricted. Left: The social welfare optimum, right: a pure Nash equilibrium.

THEOREM 3.2 (PRICE OF ANARCHY). *The price of anarchy is $|\mathcal{B}|$.*

PROOF. The price of anarchy is trivially at most $|\mathcal{B}|$ since at least one of the bakers is covered by a miller in equilibrium. For a matching lower bound, we construct the following instance: Let the set of locations be $\mathcal{L} = \{x, \ell_1, \dots, \ell_{|\mathcal{B}|}\}$ and let there be one miller. For each baker $b_i \in \{b_1, \dots, b_{|\mathcal{B}|}\}$, let the set of permissible locations be $L(b_i) = \{\ell_i, x\}$. There is a pure Nash equilibrium where each baker b_i is placed on location ℓ_i and the miller chooses ℓ_1 with only one baker covered. The social optimum has all agents on location x . See Figure 6. Thus, we have $\text{PoA} \geq |\mathcal{B}|$. \square

In contrast to the low welfare of the worst equilibrium, the best equilibrium is surprisingly good.

THEOREM 3.3 (PRICE OF STABILITY). *The PoS is $1 + \frac{\min(|\mathcal{L}|, |\mathcal{M}|-1)}{|\mathcal{M}|}$.*

The theorem follows immediately from the following two lemmas. For the upper bound, we use the strategy profile with the highest coverage and construct a pure Nash equilibrium where the millers use a subset of the miller location of this optimal state.

LEMMA 3.4 (PRICE OF STABILITY UPPER BOUND). *There exists a pure Nash equilibrium that approximates the optimal coverage by a factor of $1 + \frac{\min(|\mathcal{L}|, |\mathcal{M}|-1)}{|\mathcal{M}|}$.*

PROOF. Let L_{opt} denote the $q = \min(|\mathcal{L}|, |\mathcal{M}|)$ locations that maximize the number of covered bakers, i.e., bakers b with $L(b) \cap L_{\text{opt}} \neq \emptyset$. Clearly, the state OPT where one miller is placed on each of the locations in L_{opt} , while all bakers are assigned to a location in L_{opt} if possible, has the optimal coverage.

To obtain an equilibrium state NE with the claimed coverage, we run Algorithm 1 on a modified instance in which we remove all locations that are not in L_{opt} . To complete NE, we add the previously removed locations and assign bakers that have no available location in L_{opt} to arbitrary permissible locations.

We show that NE is still a pure Nash equilibrium: Bakers outside L_{opt} have no possible location available to them inside L_{opt} . By Theorem 2.1, bakers inside L_{opt} have no profitable deviating move to another location inside L_{opt} . Furthermore, Bakers cannot profit from moving outside L_{opt} , because there are no millers there. By Theorem 2.1, also the millers have no profitable deviating move to a location inside L_{opt} . It remains to show that millers do not want to move outside L_{opt} : Note that in NE, every baker that can be assigned to a location in L_{opt} is assigned to a location in L_{opt} . Thus, each of the $|\mathcal{M}|$ locations inside L_{opt} has at least as many bakers assigned to it as the most profitable location ℓ outside L_{opt} . Otherwise, we could construct a set L'_{opt} which covers more bakers than L_{opt} by swapping in ℓ . If not $L_{\text{opt}} = \mathcal{L}$, then L_{opt} has $|\mathcal{M}|$ locations and occupying a location in L_{opt} as the only miller is possible. Therefore, in NE any miller's utility is at least as large as the number of bakers on any location outside L_{opt} .

It remains to bound the social welfare. We denote by n the minimal utility that a miller receives in the equilibrium NE which implies a coverage of at least $|\mathcal{M}|n$. On the other hand, consider any location $\ell \in L_{\text{opt}}$ that is not covered by any miller. Equilibrium conditions imply that there are at most n bakers on ℓ as otherwise, the miller with minimal utility could improve. There are at most $q - 1$ uncovered locations in L_{opt} for NE, so there are at most $(q - 1)n$ bakers covered in OPT but not in NE. Thus, we have

$$\begin{aligned} \frac{W(\text{OPT})}{W(\text{NE})} &\leq \frac{W(\text{NE}) + (q - 1)n}{W(\text{NE})} = 1 + \frac{(q - 1)n}{W(\text{NE})} \\ &\leq 1 + \frac{(q - 1)n}{|\mathcal{M}|n} = 1 + \frac{\min(|\mathcal{L}|, |\mathcal{M}|) - 1}{|\mathcal{M}|}. \quad \square \end{aligned}$$

Note that, the algorithm implied by this lemma is more complex than Algorithm 1 because it relies on knowing the set L_{opt} with optimal coverage of bakers. Next, we match the upper bound with a lower bound on the quality of the best pure Nash equilibrium:

LEMMA 3.5 (PRICE OF STABILITY LOWER BOUND). *The price of stability is at least $1 + \frac{\min(|\mathcal{L}|, |\mathcal{M}|) - 1}{|\mathcal{M}|}$.*

PROOF. We construct the following instance for a given integer n : Let the set of locations be $\mathcal{L} = \{x, \ell_2, \dots, \ell_{|\mathcal{L}|}\}$. For location x we have $n|\mathcal{M}| + 1$ bakers restricted to only x . For each remaining location ℓ , we have n bakers restricted to only ℓ . The Nash equilibrium has all millers on x . The social optimum covers $q = \min(|\mathcal{L}|, |\mathcal{M}|)$ locations including x . See Figure 7. So we have

$$\begin{aligned} \text{PoS} &\geq \frac{|\mathcal{M}|n + 1 + n(q - 1)}{|\mathcal{M}|n + 1} = 1 + \frac{n(q - 1)}{|\mathcal{M}|n + 1} \\ \text{with } \lim_{n \rightarrow \infty} \left(1 + \frac{n(q - 1)}{|\mathcal{M}|n + 1} \right) &= 1 + \frac{\min(|\mathcal{L}|, |\mathcal{M}|) - 1}{|\mathcal{M}|}. \quad \square \end{aligned}$$




Figure 7: An instance with $|\mathcal{L}| = |\mathcal{M}|$ with the price of stability approaching $1 + \frac{\min(|\mathcal{L}|, |\mathcal{M}|) - 1}{|\mathcal{M}|}$ as $n \rightarrow \infty$. Left: The social welfare optimum, right: the only pure Nash equilibrium.

3.2 Complexity and Approximation of Optimal Solutions

We first show hardness of computing optimal solutions and optimal equilibria. To that end, we reduce from the MAXIMUM k -COVERAGE problem [28].

Definition 3.6 (MAXIMUM k -COVERAGE). Given an integer k and a set of sets S , find a subset $S' \subseteq S$ with $|S'| = k$ such that the number of included items $|\bigcup_{T \in S'} T|$ is maximal.

We now use the NP-hardness of MAXIMUM k -COVERAGE to show that finding the social welfare optimum and the socially optimal Nash equilibrium is NP-hard.

THEOREM 3.7 (SOCIAL WELFARE OPTIMUM). *Computing the social welfare optimum is NP-hard.*

PROOF. We reduce from MAXIMUM k -COVERAGE. Let the set of locations be $\mathcal{L} = S$ and for each item s in the base set $\bigcup_{T \in S} T$, we create a baker b with $L(b) = \{T \in S \mid s \in T\}$. We add k millers.

If there is an optimal solution S' that covers n items, then for the corresponding strategy profile (s, t) we place the millers at exactly the locations S' and let all bakers choose a location in S' if possible. The result is a coverage of $W(s, t) = n$.

For a social optimum (s, t) , we set S' to the set of locations that have at least one miller placed on them. Since (s, t) is a social optimum, all bakers that can be placed in S' are placed in S' and therefore the number of covered bakers is equal to the number of items contained in $\bigcup_{T \in S'} T$. \square

For optimal pure Nash equilibria, the reduction is more involved: Adding dummy bakers to each location ensures there is no pure Nash equilibrium with multiple millers on the same location.

THEOREM 3.8 (SOCALLY OPTIMAL EQUILIBRIUM). *Computing the pure Nash equilibrium with the best coverage is NP-hard.*

PROOF. Again, we reduce from MAXIMUM k -COVERAGE. Let the set of locations be $\mathcal{L} = S$ and for each item s in the base set $\bigcup_{T \in S} T$, we create a baker b with $L(b) = \{T \in S \mid s \in T\}$. For each location ℓ , we also add $q = |\bigcup_{T \in S} T| + 1$ bakers b with $L(b) = \{\ell\}$. We add k millers. We show that there is an optimal solution S' that covers n items if and only if there is a socially optimal pure Nash equilibrium (s, t) with coverage $W(s, t) = n + kq$.

If there is an optimal solution S' that covers n items, then for the corresponding miller profile \mathbf{t} we place the millers at exactly the locations S' . Given \mathbf{t} , we choose an arbitrary baker equilibrium \mathbf{s} . Note that in \mathbf{s} every baker that has a location available in S' is placed in S' . The state (\mathbf{s}, \mathbf{t}) is a pure Nash equilibrium because a miller that profitably deviates to a location $\ell \notin S'$ implies that there is a set S'' with more covered items. A miller cannot deviate to a location inside S' because, with two millers at the same location, both receive a utility of less than q . Thus, we have a pure Nash equilibrium with a coverage of $W(\mathbf{s}, \mathbf{t}) = n + kq$.

For a socially optimal Nash equilibrium (\mathbf{s}, \mathbf{t}) , we set S' to the set of locations that have at least one miller placed on them. In (\mathbf{s}, \mathbf{t}) no two millers are placed at the same location since that would yield a utility below q . Since (\mathbf{s}, \mathbf{t}) is a socially optimal Nash equilibrium, all bakers that have an available location in S' choose a location in S' . Therefore the number of covered bakers is equal to n , the number of items contained in $\bigcup_{T \in S'} T$ with kq extra bakers added. \square

With a similar technique as in our proof for the upper bound of the price of stability in Lemma 3.4, we show that the pure Nash equilibrium computed by our polynomial-time Algorithm 1 is a good approximation of the best Nash equilibrium in terms of social welfare. The added factor $\frac{e}{e-1}$ with Euler's number e results from the fact, that our algorithm greedily selects miller locations instead of choosing them optimally.

THEOREM 3.9 (NASH EQUILIBRIUM WELFARE APPROXIMATION). *Algorithm 1 computes a pure Nash equilibrium that approximates the optimal coverage by a factor of $\left(1 + \frac{\min(|\mathcal{L}|, |\mathcal{M}|) - 1}{|\mathcal{M}|}\right) \frac{e}{e-1}$.*

PROOF. We note that the greedy algorithm for the MAXIMUM k -COVERAGE problem has an approximation ratio of $\frac{e}{e-1}$ [29]. We observe that for $q = \min(|\mathcal{L}|, |\mathcal{M}|)$, this greedy solution picks exactly the locations ℓ_1, \dots, ℓ_q as defined by Algorithm 1 to place the millers. Furthermore, the algorithm assigns millers to a subset of those locations, i.e., ℓ_1, \dots, ℓ_x for some $x \leq q$ because the algorithm iteratively assigns millers to their best response location. Hence, it either assigns millers to an already covered location or the next uncovered location since the number of bakers on location ℓ_i is non-increasing with i .

Again, we denote by n the minimal utility that a miller receives in the equilibrium computed by Algorithm 1. Then the coverage in the equilibrium is at least $|\mathcal{M}|n$. On the other hand, consider any location $\ell \in \ell_{x+1}, \dots, \ell_q$ that is not covered by a miller. By the sorting of the locations, there are at most n bakers on ℓ . As there are at most $q - 1$ uncovered locations in ℓ_1, \dots, ℓ_q , the total number of bakers covered by the greedy solution but not Algorithm 1 is at most $(q - 1)n$. Thus, the ratio between the coverage of the greedy solution and the equilibrium computed by Algorithm 1 is at most

$$\frac{|\mathcal{M}|n + (q - 1)n}{|\mathcal{M}|n} = 1 + \frac{\min(|\mathcal{L}|, |\mathcal{M}|) - 1}{|\mathcal{M}|}.$$

Combining this with the factor $\frac{e}{e-1}$ of the greedy approximation yields the theorem. \square

4 CONCLUSION AND FUTURE WORK

Our model adds location choice to the classical setting of simple symmetric Fractional Hedonic Games and to Hedonic Diversity

Games with single-peaked utilities. This extension renders proving the existence of pure Nash equilibria much more challenging and we cope with this problem by carefully designing a three-stage constructive procedure that ensures that neither bakers nor millers want to deviate from the constructed state of the game. It is still an open question if this result can also be proven via a potential function approach.

The computed pure Nash equilibrium approximates the best possible social welfare by a factor of $\left(1 + \frac{\min(|\mathcal{L}|, |\mathcal{M}|) - 1}{|\mathcal{M}|}\right) \frac{e}{e-1}$. Note that the individual factors are tight: The price of stability is $1 + \frac{\min(|\mathcal{L}|, |\mathcal{M}|) - 1}{|\mathcal{M}|}$ and $\frac{e}{e-1}$ is the best achievable factor for selecting locations unless $P = NP$ [21]. However, it is an open problem, if the combination of both factors is also a tight upper bound.

An interesting future direction is to add weights to the agents. Our results on the price of anarchy and stability also hold in this setting with only minor modifications of the proofs. The bakers and the millers individually still play Congestion Games, but already for weighted agents on only one side, there is no potential function as can be seen by the improving response cycle in Figure 8.

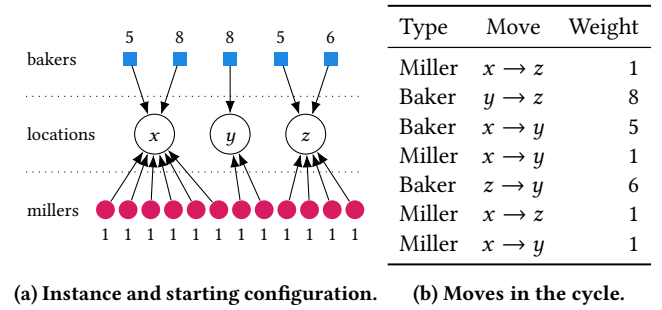


Figure 8: An instance with weighted bakers and unweighted millers that admits an improving response cycle. The edges denote the starting state locations for the cycle. All agents may choose all locations.

Another possible research direction is to restrict the millers in their location choice as well. Here, we cannot rule out the existence of a potential function. However, settling this open problem seems to be challenging, since it also remains unsolved for the related model by Gadea Harder et al. [26].

Finally, as already mentioned above, other variants of Hedonic Games can also be extended by adding a layer of location choice. We believe that investigating the impact of restricting the set of possible coalitions in such a natural way will lead to interesting new insights into coalition formation problems. Using locations, effects like geographical proximity, access restrictions, and space restrictions can be combined in simple and clean models.

ACKNOWLEDGMENTS

This work was supported by the Federal Ministry of Education and Research (BMBF) with a fellowship within the IFI program of the German Academic Exchange Service (DAAD).

REFERENCES

- [1] Aishwarya Agarwal, Edith Elkind, Jiarui Gan, Ayumi Igarashi, Warut Suksompong, and Alexandros A. Voudouris. 2021. Schelling games on graphs. *Artif. Intell.* 301 (2021), 103576. <https://doi.org/10.1016/J.ARTINT.2021.103576>
- [2] Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. 2004. The Price of Stability for Network Design with Fair Cost Allocation. In *FOCS*. 295–304. <https://doi.org/10.1109/FOCS.2004.68>
- [3] Haris Aziz, Florian Brandl, Felix Brandt, Paul Harrenstein, Martin Olsen, and Dominik Peters. 2019. Fractional Hedonic Games. *Transactions on Economics and Computation* 7, 2, Article 6 (2019), 29 pages. <https://doi.org/10.1145/3327970>
- [4] Haris Aziz, Felix Brandt, and Paul Harrenstein. 2014. Fractional hedonic games. In *AAMAS*. 5–12. <https://dl.acm.org/doi/10.5555/2615731.2615736>
- [5] Davide Bilò, Vittorio Bilò, Pascal Lenzner, and Louise Molitor. 2022. Tolerance is Necessary for Stability: Single-Peaked Swap Schelling Games. In *IJCAI*. 81–87. <https://doi.org/10.24963/ijcai.2022/12>
- [6] Vittorio Bilò, Angelo Fanelli, Michele Flammini, Gianpiero Monaco, and Luca Moscardelli. 2018. Nash stable outcomes in fractional hedonic games: Existence, efficiency and computation. *Journal of Artificial Intelligence Research* 62 (2018), 315–371. <https://doi.org/10.1613/jair.1.11211>
- [7] Duncan Black. 1948. On the rationale of group decision-making. *Journal of political economy* 56, 1 (1948), 23–34. <https://www.jstor.org/stable/1825026>
- [8] Richard Blundell. 1988. Consumer behaviour: theory and empirical evidence—a survey. *The Economic Journal* 98, 389 (1988), 16–65. <https://doi.org/10.2307/2233510>
- [9] Niclas Boehmer and Edith Elkind. 2020. Individual-Based Stability in Hedonic Diversity Games. In *AAAI*. 1822–1829. <https://doi.org/10.1609/aaai.V34i02.5549>
- [10] Anna Bogomolnaia and Matthew O. Jackson. 2002. The stability of hedonic coalition structures. *Games and Economic Behavior* 38, 2 (2002), 201–230. <https://doi.org/10.1006/game.2001.0877>
- [11] Florian Brandl, Felix Brandt, and Martin Strobel. 2015. Fractional Hedonic Games: Individual and Group Stability. In *AAMAS*. 1219–1227. <https://dl.acm.org/doi/10.5555/2772879.2773307>
- [12] Robert Brederbeck, Edith Elkind, and Ayumi Igarashi. 2019. Hedonic Diversity Games. In *AAMAS*. 565–573. <https://dl.acm.org/doi/10.5555/3306127.3331741>
- [13] Raffaello Carosi, Gianpiero Monaco, and Luca Moscardelli. 2019. Local Core Stability in Simple Symmetric Fractional Hedonic Games. In *AAMAS*. 574–582. <https://dl.acm.org/doi/10.5555/3306127.3331742>
- [14] Ankit Chauhan, Pascal Lenzner, and Louise Molitor. 2018. Schelling Segregation with Strategic Agents. In *SAGT*, Vol. 11059. 137–149. https://doi.org/10.1007/978-3-319-99660-8_13
- [15] Anthony Downs. 1957. An Economic Theory of Political Action in a Democracy. *Journal of Political Economy* 65, 2 (1957), 135–150. <https://www.jstor.org/stable/1827369>
- [16] Jacques H. Dreze and Joseph Greenberg. 1980. Hedonic coalitions: Optimality and stability. *Econometrica* (1980), 987–1003. <https://doi.org/10.2307/1912943>
- [17] Christoph Dürr and Kim Thang Nguyen. 2007. Nash Equilibria in Voronoi Games on Graphs. In *ESA*. 17–28. https://doi.org/10.1007/978-3-540-75520-3_4
- [18] Hagen Echzell, Tobias Friedrich, Pascal Lenzner, Louise Molitor, Marcus Pappik, Friedrich Schöne, Fabian Sommer, and David Stangl. 2019. Convergence and Hardness of Strategic Schelling Segregation. In *WINE*. 156–170. https://doi.org/10.1007/978-3-030-35389-6_12
- [19] Horst A. Eisel, Gilbert Laporte, and Jaques-Francois Thisse. 1993. Competitive Location Models: A Framework and Bibliography. *Transportation Science* 27, 1 (1993), 44–54. <https://doi.org/10.1287/trsc.27.1.44>
- [20] Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. 2004. The complexity of pure Nash equilibria. In *STOC*. 604–612. <https://doi.org/10.1145/1007352.1007445>
- [21] Uriel Feige. 1998. A threshold of $\ln n$ for approximating set cover. *J. ACM* 45, 4 (1998), 634–652. <https://doi.org/10.1145/285055.285059>
- [22] Rainer Feldmann, Marios Mavronicolas, and Burkhard Monien. 2009. Nash Equilibria for Voronoi Games on Transitive Graphs. In *WINE*. 280–291. https://doi.org/10.1007/978-3-642-10841-9_26
- [23] Matthias Feldotto, Pascal Lenzner, Louise Molitor, and Alexander Skopalik. 2019. From Hotelling to Load Balancing: Approximation and the Principle of Minimum Differentiation. In *AAMAS*. 1949–1951. <https://dl.acm.org/doi/10.5555/3306127.3331973>
- [24] Gaëtan Fournier and Marco Scarsini. 2019. Location Games on Networks: Existence and Efficiency of Equilibria. *Math. Oper. Res.* 44, 1 (2019), 212–235. <https://doi.org/10.1287/MOOR.2017.0921>
- [25] Tobias Friedrich, Pascal Lenzner, Louise Molitor, and Lars Seifert. 2023. Single-Peaked Jump Schelling Games. In *AAMAS*. 2899–2901. <https://dl.acm.org/doi/10.5555/3545946.3599116>
- [26] Jonathan Gadea Harder, Simon Krogmann, Pascal Lenzner, and Alexander Skopalik. 2023. Strategic Resource Selection with Homophilic Agents. In *IJCAI*. <https://doi.org/10.24963/ijcai.2023/301>
- [27] Michel X. Goemans, Li (Erran) Li, Vahab S. Mirrokni, and Marina Thottan. 2006. Market sharing games applied to content distribution in ad hoc networks. *IEEE J. Sel. Areas Commun.* 24, 5 (2006), 1020–1033. <https://doi.org/10.1109/JSAC.2006.872884>
- [28] Dorit S. Hochbaum. 1996. *Approximating covering and packing problems: set cover, vertex cover, independent set, and related problems*. PWS Publishing Co., USA, 94–143.
- [29] Dorit S. Hochbaum and Anu Pathria. 1998. Analysis of the greedy approach in problems of maximum k-coverage. *Naval Research Logistics (NRL)* 45, 6 (1998), 615–627. [https://doi.org/10.1002/\(SICI\)1520-6750\(199809\)45:6<615::AID-NAV5>3.0.CO;2-5](https://doi.org/10.1002/(SICI)1520-6750(199809)45:6<615::AID-NAV5>3.0.CO;2-5)
- [30] Harold Hotelling. 1929. Stability in Competition. *The Economic Journal* 39, 153 (1929), 41–57. <https://doi.org/10.2307/2224214>
- [31] Elon Kohlberg. 1983. Equilibrium store locations when consumers minimize travel time plus waiting time. *Economics Letters* 11, 3 (1983), 211–216. [https://doi.org/10.1016/0165-1765\(83\)90137-4](https://doi.org/10.1016/0165-1765(83)90137-4)
- [32] Elias Koutsoupias and Christos Papadimitriou. 1999. Worst-Case Equilibria. In *STACS*. 404–413. https://doi.org/10.1007/3-540-49116-3_38
- [33] Simon Krogmann, Pascal Lenzner, Louise Molitor, and Alexander Skopalik. 2021. Two-Stage Facility Location Games with Strategic Clients and Facilities. In *IJCAI*. 292–298. <https://doi.org/10.24963/ijcai.2021/41>
- [34] Simon Krogmann, Pascal Lenzner, and Alexander Skopalik. 2023. Strategic Facility Location with Clients That Minimize Total Waiting Time. In *AAAI*. 5714–5721. <https://doi.org/10.1609/aaai.V37i5.25709>
- [35] Simon Krogmann, Pascal Lenzner, Alexander Skopalik, Marc Uetz, and Marnix C. Vos. 2024. Equilibria in Two-Stage Facility Location with Atomic Clients. In *IJCAI*. 2842–2850. <https://doi.org/10.24963/ijcai.2024/315>
- [36] Gianpiero Monaco and Luca Moscardelli. 2023. Nash Stability in Fractional Hedonic Games with Bounded Size Coalitions. In *WINE*, Vol. 14413. 509–526. https://doi.org/10.1007/978-3-031-48974-7_29
- [37] Gianpiero Monaco, Luca Moscardelli, and Yllka Velaj. 2019. On the Performance of Stable Outcomes in Modified Fractional Hedonic Games with Egalitarian Social Welfare. In *AAMAS*. 873–881. <https://dl.acm.org/doi/10.5555/3306127.3331780>
- [38] Gianpiero Monaco, Luca Moscardelli, and Yllka Velaj. 2020. Stable outcomes in modified fractional hedonic games. *Autonomous Agents Multi Agent Systems* 34, 1 (2020), 4. <https://doi.org/10.1007/S10458-019-09431-Z>
- [39] Dov Monderer and Lloyd S. Shapley. 1996. Potential Games. *Games and Economic Behavior* 14, 1 (1996), 124–143. <https://doi.org/10.1006/game.1996.0044>
- [40] James B. Orlin. 1993. A faster strongly polynomial minimum cost flow algorithm. *Operations Research* 41, 2 (1993), 338–350. <https://doi.org/10.1287/opre.41.2.338>
- [41] Susan Hesse Owen and Mark S. Daskin. 1998. Strategic facility location: A review. *Eur. J. Oper. Res.* 111, 3 (1998), 423–447. [https://doi.org/10.1016/S0377-2217\(98\)00186-6](https://doi.org/10.1016/S0377-2217(98)00186-6)
- [42] Hans Peters, Marc Schröder, and Dries Vermeulen. 2018. Hotelling’s location model with negative network externalities. *Int. J. Game Theory* 47, 3 (2018), 811–837. <https://doi.org/10.1007/S00182-018-0615-0>
- [43] Charles S. ReVelle and Horst A. Eisel. 2005. Location Analysis: A Synthesis and Survey. *European Journal of Operational Research* 165, 1 (2005), 1–19. <https://doi.org/10.1016/j.ejor.2003.11.032>
- [44] Robert W. Rosenthal. 1973. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2 (1973), 65–67. <https://doi.org/10.1007/BF01737559>
- [45] Leon Schiffrman, Aron O’Cass, Angela Paladino, and Jamie Carlson. 2013. *Consumer behaviour*. Pearson Higher Education AU.
- [46] Daniel Schmand, Marc Schröder, and Alexander Skopalik. 2019. Network Investment Games with Wardrop Followers. In *ICALP*. 151:1–151:14. <https://doi.org/10.4230/LIPICS.ICALP.2019.151>
- [47] Adrian Vetta. 2002. Nash Equilibria in Competitive Societies, with Applications to Facility Location, Traffic Routing and Auctions. In *FOCS*. 416–425. <https://doi.org/10.1109/SFCS.2002.1181966>