

Schelling Models with Localized Social Influence: A Game-Theoretic Framework

Hau Chan*
University of Nebraska-Lincoln
Lincoln, Nebraska
hchan3@unl.edu

Mohammad T. Irfan*
Bowdoin College
Brunswick, Maine
mirfan@bowdoin.edu

Cuong Viet Than
University of Nebraska-Lincoln
Lincoln, Nebraska
cthan2@huskers.unl.edu

ABSTRACT

We propose a game-theoretic approach to generalizing the classical Schelling model. At the core of our model are two features that did not receive much attention before. First, we allow multiple individuals to occupy the same location. Second, each individual's choice of location is influenced by their social network neighbors that also choose the same location. In addition, an individual's choice is influenced by others in the adjacent locations in a network-structured way, which captures the main spirit of the classical Schelling model and its numerous extensions. Our solution concept is a stable configuration represented as a pure-strategy Nash equilibrium (PSNE). We show that even for various special cases of the problem, computing or counting PSNE is provably hard. We give algorithms for computing PSNE, including efficient algorithms for several special cases. We highlight some of the attractive features of our model, such as predicting very few PSNE, through experiments.

KEYWORDS

Schelling model; social networks; social influence; computational game theory; Nash equilibrium; propagation algorithms

ACM Reference Format:

Hau Chan, Mohammad T. Irfan, and Cuong Viet Than. 2020. Schelling Models with Localized Social Influence: A Game-Theoretic Framework. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*, Auckland, New Zealand, May 9–13, 2020, IFAAMAS, 9 pages.

1 INTRODUCTION

Residential segregation by race has been historically well documented in areas with diverse populations. Even the most recent studies show that urban living is segregated with respect to race, ethnicity, and social status. For example, the 2010 U.S. census data gives clear evidence of segregation by race in major metropolitan areas like Chicago, Washington, D.C., and Houston.¹ The 2011 U.K. census data also shows a similar pattern in major cities.²

Although early work in mathematical social sciences focused on how individual choices lead to a segregated collective outcomes [8, 23–25], recent studies driven by census data show broader implications of segregation. For example, the Chicago Reader reports

findings by the Social Impact Research Center that economic opportunities in Chicago are very much correlated with the racial composition of a community.³ There are examples of Chicago neighborhoods with over 90% black population where the unemployment rates are about nine times that of some of the majority white neighborhoods. To effectively address issues like poverty and violence that are intertwined with segregation, there is a need for more sophisticated models that generalize the classical Schelling model [23–25] and its extensions. This paper presents work in this direction.

Any study of segregation must begin with the seminal work by Noble prize-winning economist Thomas Schelling [24, 25]. Aimed with the goal of understanding and modeling the process of segregation through the lens of individual choices, Schelling introduced a dynamic model of segregation over time with two *types* of individuals. Here, types may represent race or other homophily criteria [8]. The model starts with initial locations of individuals within set of stylized locations like a grid. An individual i is happy or satisfied if at least t_i fraction of i 's neighbors are of the same type as i . At each time-step of the dynamic process, unsatisfied individuals from the previous time step move to different locations (via an algorithm or some random process) where they can be satisfied. The model enforces the rule that no two individuals can be in the same location at any time. Furthermore, individual i 's threshold t_i should not be interpreted as i 's penchant to move to a neighborhood where i can be among the majority. In contrast, t_i in Schelling's model has the connotation of i 's desire to avoid being an extreme minority. Schelling's model shows that segregation happens even when individuals do not desire to be in a majority neighborhood.

Although there has been a continuous stream of multidisciplinary research on Schelling's model [7, 12, 20, 27], the topic only started gaining traction within computer science fairly recently [1–3, 10, 13]. To our knowledge, Chauhan et al. were the first to give a strictly game-theoretic model of segregation where individuals have preferences over networked locations and choose a location strategically [3]. Their focus is on the convergence properties when the locations are connected in the form of a ring or a regular graph. In this paper, we call the graph connecting locations the *location graph*. Chauhan et al.'s consideration of very specific types of location graphs can be attributed to the complexity of their model. They allow thresholds in the spirit of the classical Schelling models. They also allow individual preferences over locations.

Soon afterwards, Elkind et al. gave a static game-theoretic model of the Schelling segregation (a.k.a. Schelling Games) that relaxes Chauhan et al.'s model by completely getting rid of thresholds and to some extent also getting rid of location preferences [10]. Given a

*HC and MTI contributed equally and are co-first authors.

¹<https://www.washingtonpost.com/graphics/2018/national/segregation-us-cities/>

²<https://www.opendemocracy.net/en/is-segregation-on-increase-in-uk/>

³<https://www.chicagoreader.com/chicago/still-separate-unequal-and-ignored/Content?oid=16347785>

location graph, every *strategic* agent in their model simultaneously and strategically chooses where to live based on a location-wise quantification of where they can be “happiest.” Given a strategy profile, the utility of agent i is defined to be the fraction of i ’s neighbors in the location graph that are of the same type as i over the total number of neighbors of i . Our paper generalizes Elkind et al.’s model in several ways. We allow multiple occupants in a location and consider capacity constraints on the locations, whereas Elkind et al. [10], Chauhan et al. [3], and not surprisingly, Schelling’s original models [24, 25] allowed at most one agent in each location. More importantly, we consider a weighted and directed social network among the agents, where an individual can be influenced by others in varying magnitude and polarity. This social network component got brief attention in Elkind et al.’s work, where they modeled the social network as an unweighted graph [10]. However, its full exploration has remained open, especially when the social network is a directed, weighted graph. This is one of our main goals.

Very recently, Echzell et al. presented some very interesting results on the convergence of best response dynamics in Schelling games [9]. They showed “knife-edge” properties of threshold values between convergence and non-convergence. In this paper, however, we do not deal with dynamics.

In sum, in previous models, the happiness or utility of the agents is defined to be some (weighted) cardinal values of the types of other agents living in the surrounding areas. We call this the *location effect*, which is rooted in the classical Schelling models. However, agents living in the same location did not get any attention, primarily because the previous models implicitly set a capacity of 1 for each location. We address this by allowing multiple occupants in each location and accounting for the influence that an agent’s social network neighbors living in the same location have on that agent. We call this the *localized social influence*. This is one of our major conceptual contributions in this paper. We address the following fundamental question that did not receive any attention before.

*How should we model the utility of an agent when in addition to the **location effect**, the agent takes into account the **localized social influence** among those agents that are living in the same location?*

We argue that when an agent makes a strategic decision like choosing where to live, the agent’s choice depends on other agents in their social circles or networks. Several research has illustrated such phenomena, ranging from health and behavioral choices [5, 6, 11] to voting [14–16] to economic decision making [19, 26]. Roughly speaking, an agent tends to make the same decision as their peers who have the most influence on them. In our context, if many of the agent’s influential peers live in a particular location, the agent would also have a lot of incentive to live in the same location. Thus, when an agent makes a decision on where to live, the agent is influenced (with varying influence levels) by the decisions of other agents in her social circle. We view this mutual interdependency among the agents in a game-theoretic way. Motivating examples of this framework can be found in people’s choices of school districts, parks, and other shared public spaces, where both social influence and location effect play a role.⁴ Fig. 1 provides an illustration.

⁴We sincerely thank an anonymous reviewer for suggesting these examples.

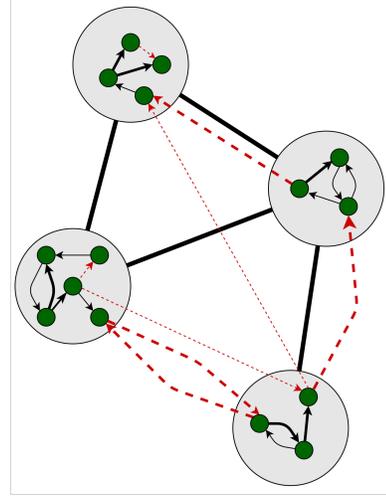


Figure 1: Illustration of our model: There are two graphs. The first one is the location graph with four location nodes (big, gray nodes) and the black undirected edges among them. The location graph is inspired by the classical Schelling model and its recent extensions [3, 10]. The second graph represents the social network among the agents. It consists of the green nodes (much smaller in size compared to the location nodes) and red and black directed edges. The black solid edges represent positive influence. The red dashed edges represent negative influence. The thickness of the directed edges stands for the magnitude of influence. The placement of the agents inside the location nodes signifies the choices made by the agents. Allowing multiple agents within a location and accounting for the localized social influence among the agents sharing the same location are novel contributions of this paper.

Table 1: Summary of Results

Problem in SG-LSI	Complexity
Is there a PSNE?	NP-complete
Is there a socially optimal PSNE?	NP-complete
Is there a locally optimal PSNE?	NP-complete
How many PSNE?	#P-complete
Finding a PSNE in 2-location tree SG-LSI	$O(n\Delta)$
Finding a PSNE in weighted-structured SG-LSI	$O(nm)$

To model localized social influence in our setting, we use a class of succinctly representable graphical game of parametric form called *influence games* [15, 16]. Influence games can represent varying positive and negative influence weights among individuals using potentially asymmetric edges. They also allow varying levels of tolerance to influence (or threshold values) across individuals. Unlike the widely studied influence maximization problem [4, 18, 21], influence games model collective outcomes as Nash equilibria in a strictly game-theoretic fashion. This not only makes influence

games attractive to our case but also has potential to broaden recent computational works on Schelling games [1, 3, 10].

Contributions. In this paper, we study the game-theoretic modeling and computational aspects of (1) the location effect as studied in the classical Schelling model and its extensions, and (2) the localized social influence, which we introduce here. In our model, agents make simultaneous and strategic decisions of selecting a location to live. We first introduce a general game-theoretic framework to capture (1) and (2). The framework encompasses the existing Schelling Games [10] and Influence Games [16] to model the location effect in Schelling settings and the localized social influence to capture network effects, respectively. By virtue of modeling localized social influence, unlike the previous models, our framework allows more than one agent to occupy the same location simultaneously, subject to capacity constraints. We show the hardness of the problem for multiple variants. We design algorithms for special cases like a tree-structured social network with two locations. For this tree case, we give a polynomial-time algorithm, whereas the standard TreeNash algorithm runs in exponential time [17]. Through experiments, we show that incorporating localized social influence leads to a more predictive model by reducing the number of PSNE.

Organization. In Section 2, we present our model, including the two features that differentiates our model from previous ones – localized social influence and allowing multiple agents within the same location. We also connect our study to various prior studies at a technical level. In Section 3, we establish hardness results for various special cases of our problem. This clearly shows that the computational problem in its general form is intractable unless $P = NP$. Section 4 deals with algorithmic results, where we focus on special cases of the problem due to the hardness of the general case. In Section 5, we present some interesting experimental results that clearly show the value of modeling localized social influence. It should be noted here that experimental work did not get much attention in recent research within computer science [2, 3, 10, 13]. Section 6 concludes the paper by outlining some open problems on the inefficiency of equilibria.

2 A GAME-THEORETIC FRAMEWORK

In this section, we introduce our game-theoretic model of Schelling Games with Localized Social Influence (SG-LSI). As discussed earlier, SG-LSI consists of two major components: *localized social influence* and *location effect*. The localized social influence component describes how an agent will be affected by other agents in his/her social circles living in the same location. The location effect component measures how an agent will be affected by other agents living in different locations. We begin by discussing each of these components. We then discuss how we can combine the components together to define the utility functions of the agents.

Localized Social Influence. To model localized social influence, we use a social network in the form a weighted, directed graph to define the relationship strengths or ties between any pair of agents. In particular, let $N = \{1, \dots, n\}$ be the set of n agents. Let $G = (N, E, w)$ be a directed social network or graph with the edge set E and the weight function $w : N \times N \rightarrow \mathbb{R}$ specifying the

“influence” weights. That is, $w(i, j) \neq 0$ if and only if $(i, j) \in E$.⁵ We allow the weights to be arbitrarily positive, negative, or zero. A positive weight from i to j (i.e., $w(i, j) > 0$) indicates a positive influence from i to j . That is, if agent i partakes a certain action, then agent j is also influenced to partake the same action. On the other hand, if $w(i, j) < 0$, then agent j would have a less desire to partake the same action of i . In our setting, an agent should have more incentive to select a location in which many of its positively influenced neighbors in G have also selected the same location. On the other hand, agent i would want to stay away from selecting the locations of those of negatively influencing neighbors in G . We use $\mathcal{P}(i) = \{j \in N : w(j, i) \neq 0\}$ to denote the set of parents of i .

Location Effect. In the standard Schelling setting, each agent selects a location from a set of possible locations. Let $L = \{1, \dots, m\}$ be the set of m possible locations. The locations can be connected via some (undirected) graph structure $G^L = (L, E^L)$ where $j, j' \in L$ are adjacent to each other if and only if $\{j, j'\} \in E^L$. For each location $j \in L$, there is a capacity $c_j \geq 1$ specifying the maximum number of agents that can live in location $j \in L$. In the standard Schelling model, $c_j = 1$ for all location $j \in L$. The location effect on an agent in the standard Schelling setting is defined to be some cardinal value of other agents in the neighboring locations of the agent’s selected location. For a location $j \in L$, we define $\mathcal{N}(j) = \{j' \in L : \{j, j'\} \in E^L\}$ to be the set of j ’s neighboring locations. We assume that the total capacity can accommodate all of the n agents in the system (i.e., $\sum_{j \in L} c_j \geq n$).

Schelling Games with Localized Social Influence (SG-LSI). Now that we have defined the two main components of SG-LSI, we are ready to define the SG-LSI formally. We represent an SG-LSI with the tuple $\mathcal{G} = (G, G^L, \{c_j\}_{j \in L}, S, \{\alpha_i\}_{i \in N}, \{\lambda_i\}_{i \in N}, \{b_{i,j}\}_{i \in N, j \in L}, f, \{u_i\}_{i \in N})$, where the influence network G , location network G^L , and capacities c_j are defined above. We next define rest of the terms. Let $S = \{1, \dots, m\}$ be the action/strategy set of each agent $i \in N$. In other words, the set of actions for any agent is the same as the set of locations L . We let S^n to be the set of (pure) action profiles and $\mathbf{a} = (a_1, \dots, a_n) \in S^n$ be a (pure) action profile. We denote (a_i, \mathbf{a}_{-i}) (or $(a_i, \mathbf{a}_{\mathcal{P}(i)})$) to specify the action of i given the action profile of other agents beside i (or the parents of i). Let $A(\mathbf{a}, j) = \{i \in N : a_i = j\}$ be the set of agents selecting location j . Given an action profile \mathbf{a} , agent i ’s utility is defined as

$$u_i(a_i, \mathbf{a}_{-i}) = \begin{cases} -\infty & \text{if } |A(\mathbf{a}, a_i)| > c_{a_i} \\ \alpha_i(\sum_{k \in A(\mathbf{a}, a_i)} w(k, i) + b_{i, a_i}) + \lambda_i f_i(\mathcal{N}(a_i), \mathbf{a}) & \text{if } |A(\mathbf{a}, a_i)| \leq c_{a_i}. \end{cases} \quad (1)$$

When agent i ’s selected location cannot accommodate all of the agents that also select i ’s location (i.e., $|A(\mathbf{a}, a_i)| > c_{a_i}$), the utility of agent i is extremely negative.⁶ When the i ’s selected location can accommodate all of the agents that select i ’s location (i.e.,

⁵We use a tuple when referring to a directed edge and an unordered set when referring to an undirected edge.

⁶The utility in this case is consistent with the standard Schelling model [24, 25] and Schelling games [10] when more than one agent selects the same location. We could also define the utility function based on some tie-breaking ordering to determine which of the $|A(\mathbf{a}, a_i)| > c_{a_i}$ agents the location a_i will accommodate.

$|A(\mathbf{a}, a_i)| \leq c_{a_i}$), the utility of agent i depends on *localized social influence* and *location effect*.

The localized social influence term for i is defined to be the sum of the influences from i 's parents who select the same location as i and i 's intrinsic threshold, b_{i,a_i} , for a location a_i . This threshold term models an agent's location preferences, which is an important feature of our model.

The location effect term is defined to be some computable function f_i , which is a function of agents that select the adjacent locations of the agent's location under the location graph G^L . Similar location effect term is considered in the classical Schelling model as well as almost all variants of it [3, 10, 24, 25].

The $\alpha_i \geq 0$ and $\lambda_i \geq 0$ terms serve as a trade off between localized social influence and location effect.

Given an instance \mathcal{G} of SG-LSI, we are interested in the question of computing a pure-strategy Nash equilibrium (PSNE).

Definition 2.1. Given any SG-LSI instance \mathcal{G} , a pure-strategy profile $\mathbf{a}^* \in S^n$ is a pure-strategy Nash equilibrium (PSNE) of \mathcal{G} if and only if $u_i(a_i^*, \mathbf{a}_{-i}^*) \geq u_i(a_i, \mathbf{a}_{-i}^*)$ for all $i \in N$ and $a_i \in S$.

Connection to Schelling Games. Below, we show how we can transform the SG-LSI to the Schelling Game (SG) [10]. In a SG, there is an undirected location graph $L = (V, E)$. The agents are divided into multiple battling factions (or types), where agents of the same type are friends and different types are enemies. Furthermore, the agents are partitioned into stubborn agents S and strategic agent R where the stubborn agents always want to select some fixed locations. Given a feasible assignment vector (or action profile in our terminology) $\mathbf{v} = (v_1, \dots, v_n) \in V^n$, the utility of agent i is $\hat{u}_i(v_i, \mathbf{v}_{-i}) = \frac{f_i(\mathbf{v})}{f_i(\mathbf{v}) + e_i(\mathbf{v})}$, where $f_i(\mathbf{v})$ is the number of i 's friends that select some adjacent location of i in the location graph G and $e_i(\mathbf{v})$ is the number of enemies of i that select some adjacent location of i . If $f_i(\mathbf{v}) = 0$, then $\hat{u}_i(v_i, \mathbf{v}_{-i}) = 0$. (Note that we use the term f_i to denote the location effect in our model, whereas Elkind et al. use f_i to count the number of friends of i in adjacent locations. Furthermore, we use L to denote the set of locations, not the location graph.)

It is easy to see that we can construct an SG-LSI instance with n agents for any instance of SG with n strategic agents (we do not model stubborn agents). In particular, consider

$$\mathcal{G} = (G, G^L, S, \{c_j\}_{j \in L}, \{\alpha_i\}_{i \in N}, \{\lambda_i\}_{i \in N}, \{b_{i,j}\}_{i \in N, j \in L}, f, \{u_i\}_{i \in N}),$$

where $G^L = (L, E^L)$, $c_j = 1$, $f = \hat{u}_i$, $\lambda_i = 1$, and $\alpha_i = 0$ for all $i \in N$ and $j \in L$. The types of the agents can be defined with respect to G .

Connection to Linear Influence Games. We now establish a connection between SG-LSI and Linear Influence games (LIG) [15, 16]. LIG is used to model the adoption behavior among strategic agents in complex social networks where agents influence each other in varying magnitudes and polarities. In an LIG, we have n agents. Agent i 's action is denoted by $x_i \in \{-1, 1\}$. The influence function of each individual i is defined as $f_i(\mathbf{x}_{-i}) = \sum_{j \neq i} w_{ji} x_j - b_i$ where for any other individual j , $w_{ji} \in \mathbb{R}$ is a *weight* parameter quantifying the "influence factor" that j has on i , and $b_i \in \mathbb{R}$ is a threshold parameter for i 's level of "tolerance" for negative effect. The utility function $u_i : \{-1, 1\}^n \rightarrow \mathbb{R}$ as $u_i(x_i, \mathbf{x}_{-i}) = x_i f_i(\mathbf{x}_{-i})$, where \mathbf{x}_{-i}

denotes the joint-action of all players except i . (Note that we use f_i to denote location effect.)

It is easy to see that we can construct an SG-LSI instance with n agent for any instance of LIG with n agents. In particular, consider

$$\mathcal{G} = (G, G^L, S, \{c_j\}_{j \in L}, \{\alpha_i\}_{i \in N}, \{\lambda_i\}_{i \in N}, \{b_{i,j}\}_{i \in N, j \in L}, f, \{u_i\}_{i \in N}),$$

where $G^L = (L = \{1, 2\}, E^L)$, $c_j = n$, $\lambda_i = 0$, and $\alpha_i = 1$ for all $i \in N$ and $j \in L$. Since $\lambda_i = 0$, the location graph and the function f can be defined arbitrarily. The utility function of each player i in \mathcal{G} can be defined accordingly (see Section 3).

3 COMPUTATIONAL COMPLEXITY OF COMPUTING A PSNE OF SG-LSI

In this section, we discuss the complexity of computing a PSNE of an SG-LSI. We will show that computing a PSNE is NP-complete and counting the number of PSNE is #P-complete via a reduction from LIG. We will also show that computing a PSNE that maximizes some social welfare measure is NP-hard.

3.1 Computing a General PSNE of SG-LSI

In this section, we prove the following hardness results.

THEOREM 3.1. *It is NP-complete to decide whether there exists a PSNE in an SG-LSI even with two locations, no location effect, and a bipartite social network.*

THEOREM 3.2. *It is #P-complete to count the number of PSNE in an SG-LSI even with two locations, no location effect, and a bipartite or star social network.*

Let \mathcal{LG} be any n -agent $\{-1, 1\}$ -action LIG with the influence weights \mathbf{w} and threshold values \mathbf{b} . We will slightly abuse the notation here by treating influence weight w_{ji} in LIG the same as $w(j, i)$ in SG-LSI. We show a polynomial-time reduction from any LIG instance \mathcal{LG} to an instance \mathcal{G} of SG-LSI such that there exists a PSNE in \mathcal{LG} if and only if there exists a PSNE in \mathcal{G} . The following two definitions are in the context of an LIG.

Definition 3.3 (Best-Response Correspondence [16]). Given $\mathbf{x}_{-i} \in \{-1, 1\}^{n-1}$, the *best-response correspondence* $\mathcal{BR}_i^{\mathcal{LG}} : \{-1, 1\}^{n-1} \rightarrow 2^{\{-1, 1\}}$ of a player i of an LIG \mathcal{LG} is defined as follows.

$$\mathcal{BR}_i^{\mathcal{LG}}(\mathbf{x}_{-i}) \equiv \arg \max_{x_i \in \{-1, 1\}} u_i(x_i, \mathbf{x}_{-i}).$$

Definition 3.4 (Pure-Strategy Nash Equilibrium [16]). A *pure-strategy Nash equilibrium (PSNE)* of an LIG \mathcal{LG} is an action assignment $\mathbf{x}^* \in \{-1, 1\}^n$ that satisfies the following condition. Every agent i 's action x_i^* is a simultaneous best-response to the actions \mathbf{x}_{-i}^* of the rest.

The utility of agent i of LIG \mathcal{LG} is $u_i(x_i, \mathbf{x}_{-i}) = x_i (\sum_j w_{ji} x_j - b_i)$. Here, we assume $w_{ii} = 0$ for all i . This does not change the above hardness results on LIGs. Let $\mathbf{x}^* \in \{-1, 1\}^n$ be a PSNE of \mathcal{LG} . It must be the case that for every agent i ,

$$x_i^* \left(\sum_j w_{ji} x_j^* - b_i \right) \geq 0. \quad (2)$$

Otherwise, i would have incentive to unilaterally switch to $-x_i^*$.

Given \mathcal{LG} , we construct an instance \mathcal{G} of SG-LSI as follows. First, let $\alpha_i = 1$ and $\lambda_i = 0$ for all agents i . We have two locations 1 and 2 in \mathcal{G} , each with a capacity of n . The agents in \mathcal{G} and \mathcal{LG} are the same. The network structure and the weight function in \mathcal{G} are also the same as in \mathcal{LG} . Corresponding to each agent i 's threshold b_i in \mathcal{LG} , we define thresholds $b_{i,1} = \frac{-b_i}{2}$ and $b_{i,2} = \frac{b_i}{2}$. Finally, we assume that the actions 1 and -1 in \mathcal{LG} correspond to the locations 1 and 2 in \mathcal{G} .

LEMMA 3.5. *There exists a PSNE in the LIG instance \mathcal{LG} if and only if there exists a PSNE in the SG-LSI \mathcal{G} .*

PROOF. Let $\mathbf{x}^* \in \{-1, 1\}^n$ be a PSNE of \mathcal{LG} and let $\mathbf{a}^* \in \{1, 2\}^n$ be the corresponding joint action in \mathcal{G} , where $a_i^* = 1$ if and only if $x_i^* = 1$. Using (2), we obtain the following sequence of equivalences. Here, $\mathbb{1}[\cdot]$ stands for the boolean indicator function.

$$\begin{aligned} x_i^* \left(\sum_j w_{ji} x_j^* - b_i \right) &\geq 0. \\ \sum_j w_{ji} \frac{1 + x_i^* x_j^*}{2} - \sum_j w_{ji} \frac{1 - x_i^* x_j^*}{2} - x_i^* \left(\frac{b_i}{2} - \frac{-b_i}{2} \right) &\geq 0. \\ \sum_j w_{ji} \mathbb{1}[x_i^* = x_j^*] - \sum_j w_{ji} \mathbb{1}[x_i^* \neq x_j^*] + x_i^* \left(\frac{-b_i}{2} - \frac{b_i}{2} \right) &\geq 0. \end{aligned} \quad (3)$$

We will now translate the PSNE \mathbf{x}^* of \mathcal{LG} to the joint action \mathbf{a}^* of \mathcal{G} and show that \mathbf{a}^* is a PSNE of \mathcal{G} . For this, note that $x_i^* \left(\frac{-b_i}{2} - \frac{b_i}{2} \right)$ translates to $\frac{-b_i}{2} - \frac{b_i}{2} = b_{i,1} - b_{i,2}$ when $x_i^* = 1$ (i.e., $a_i^* = 1$). It translates to $b_{i,2} - b_{i,1}$ when $x_i^* = -1$ (i.e., $a_i^* = 2$). Therefore, in general, $x_i^* \left(\frac{-b_i}{2} - \frac{b_i}{2} \right)$ translates to $b_{i,a_i^*} - b_{i,\bar{a}_i^*}$ where \bar{a}_i^* denotes the complement of the action a_i^* . We obtain from (3):

$$\begin{aligned} \sum_j w_{ji} \mathbb{1}[a_i^* = a_j^*] - \sum_j w_{ji} \mathbb{1}[a_i^* \neq a_j^*] + b_{i,a_i^*} - b_{i,\bar{a}_i^*} &\geq 0. \\ \sum_j w_{ji} \mathbb{1}[a_i^* = a_j^*] + b_{i,a_i^*} &\geq \sum_j w_{ji} \mathbb{1}[a_i^* \neq a_j^*] + b_{i,\bar{a}_i^*}. \\ \sum_{j \in S_{\mathbf{a}^*}(a_i^*)} w_{ji} + b_{i,a_i^*} &\geq \sum_{j \in S_{\mathbf{a}^*}(\bar{a}_i^*)} w_{ji} + b_{i,\bar{a}_i^*}. \end{aligned}$$

Therefore, \mathbf{a}^* is a PSNE of the SG-LSI instance \mathcal{G} . We can reverse the above argument and show that if \mathbf{a}^* is a PSNE of \mathcal{G} then the corresponding joint action \mathbf{x}^* is a PSNE of \mathcal{LG} . \square

Lemma 3.5 leads us to Theorems 3.1 and 3.2 by noting that it is NP-complete to decide whether there exists a PSNE and #P-complete to count the number of PSNE in an LIG [16].

3.2 Computing a Specific PSNE of SG-LSI

In this section, we show that the problem of computing a PSNE of SG-LSI that maximizes social welfare is NP-hard. The social welfare of a PSNE $\mathbf{a}^* \in S^n$ is defined to be $SW(\mathbf{a}^*) = \sum_{i \in N} u_i(\mathbf{a}^*)$. Thus, we want to find $\mathbf{a}_{opt}^* \in \arg \max_{\mathbf{a}^* \text{ is a PSNE}} SW(\mathbf{a}^*)$.

To show that the problem is NP-hard, we reduce from the k -colorable problem. The k -colorable problem is known to be NP-hard for $k \geq 3$. In an instance of k -colorable problem $C = (G, k)$, we are giving an undirected graph $G = (V, E)$ and an integer $k \geq 3$. We

want to know if the graph G is k -colorable. That is, is there exist a way to color the vertex of G with k colors so that no two adjacent vertices have the same color. More formally, is there a function $\chi : V \rightarrow \{1, \dots, k\}$ such that $\chi(v) \neq \chi(v')$ for any $\{v, v'\} \in E$.

Given an instance of $C = (G, V)$ with n nodes, we reduce it to an instance of SG-LSI with n agents. The social graph in the reduced SG-LSI is the same undirected graph, but the weight function $w(v, v') < 0$ is strictly negative for any $\{v, v'\} \in E$. The number of locations is set to k where $L = \{1, \dots, k\}$. The $\alpha_i = 1$, $\lambda_i = 0$, and $b_{i,j} = 0$ for any $i \in N$ and $j \in L$. Since $\lambda_i = 0$, the structure of the locations and the f function can be defined arbitrarily. Finally, $c_j = n$ for any $j \in L$. In the following, we show that C is k -colorable if and only if there a PSNE \mathbf{a}^* with $SW(\mathbf{a}^*) \geq 0$.

LEMMA 3.6. *The k -colorable problem instance C is k -colorable if and only if there is a PSNE \mathbf{a}^* in the corresponding SG-LSI instance \mathcal{G} such that $SW(\mathbf{a}^*) \geq 0$.*

PROOF. We first show that if C is k -colorable, then the k -colorable solution can be used to construct a PSNE \mathbf{a}^* with $SW(\mathbf{a}^*) \geq 0$ in \mathcal{G} .

Let χ be a solution that maps V to one of the k colors. We create a strategy profile $\mathbf{a}^* = (a_1, \dots, a_n)$ where $a_i = \chi(i)$ for each $i \in V$ in the corresponding graph. Because there is no pair of two vertices that have same color (i.e., $\chi(v) \neq \chi(v')$ for any $\{v, v'\} \in E$), the agents are in different locations than their neighbors, and the capacity of location j is always large enough for any subset of agents, the utility of agent i is $u_i(a_i^*, \mathbf{a}_{-i}^*) = \sum_{t \in A(\mathbf{a}^*, a_i^*)} w(t, i) = 0$. Since the influence weights are negative $u_i(a_i^*, \mathbf{a}_{-i}^*) \geq u_i(a_i, \mathbf{a}_{-i}^*)$ for any $a_i \in S$ and $i \in N$. As a result, \mathbf{a}^* is a PSNE and $SW(\mathbf{a}^*) \geq 0$.

Now, suppose that we have a PSNE \mathbf{a}^* with $SW(\mathbf{a}^*) \geq 0$. It is not hard to see that $u_i(a_i^*, \mathbf{a}_{-i}^*) = 0$ for any $i \in N$ since all of the edges in the graph have negative edges and $u_i(a_i^*, \mathbf{a}_{-i}^*) \leq 0$. Thus, in any PSNE \mathbf{a}^* with $SW(\mathbf{a}^*) \geq 0$, $a_i \neq a_{i'}$ for any $w(i, i') < 0$. We can construct a k -colorable solution for C by defining $\chi(i) = a_i$ for each $i \in N$.

Using the above two arguments, we have proved our lemma. \square

From Lemma 3.6, we have the following.

THEOREM 3.7. *It is NP-hard to compute a PSNE that maximizes the social welfare in an SG-LSI even on social network with only negative edges and no location effect for more than three locations.*

Our next question concerns with computing a PSNE that maximizes the social welfare of a particular location. Let $SW(\mathbf{a}^*) = \sum_{j \in L} SW_j(\mathbf{a}^*) = \sum_{j \in L} \sum_{i \in A(\mathbf{a}^*, j)} u_i(a_i^*, \mathbf{a}_{-i}^*)$ where $SW_j(\mathbf{a}^*)$ is defined to be the social welfare of location $j \in L$ (which is the sum of the utilities of the agents in location j). For a given location $j \in L$, we want to find $\mathbf{a}_{opt,j}^* \in \arg \max_{\mathbf{a}^* \text{ is a PSNE}} SW_j(\mathbf{a}^*)$. Theorem 3.7 shows it is NP-hard to maximize social welfare of a particular location with only negative edges. Our below result complementaries the above result by showing it is still hold for positive edges.

THEOREM 3.8. *It is NP-hard to compute a PSNE that maximizes the social welfare of a location in an SG-LSI even on a social network with positive unit-weighted edges and no location effect.*

PROOF. We prove the claim by reducing from the k -clique problem, which is known to be NP-complete. In the k -clique problem,

we are given a graph $G = (V, E)$ and an integer k . We want to know if there is a clique (or a complete graph) of size k in G .

Given a k -clique instance C , we can construct an SG-LSI instance \mathcal{G} with the same social network as G of uniform positive weight of 1 (i.e., $w(i, i') = 1$ for $\{i, i'\} \in E$). The number of locations is set to $n - k + 1$ where $L = \{1, \dots, n - k + 1\}$. The $\alpha_i = 1$, $\lambda_i = 0$, and $b_{i,j} = 0$ for any $i \in N$ and $j \in L$. Since $\lambda_i = 0$, the structure of the locations and the f function can be defined arbitrarily. Finally, $c_1 = k$ and $c_j = 1$ for $j \in L \setminus \{1\}$. Our goal is to show that C has a clique of size k if and only if there is a PSNE \mathbf{a}^* such that $SW_1(\mathbf{a}^*) \geq k(k - 1)$.

It is not hard to see that if there is a k clique, we can place all of the k agents into location 1 (which has the capacity of k) and the other agents in the other locations with capacity 1. The agents have no incentive to deviate (i.e., \mathbf{a}^* is a PSNE) and $SW_1(\mathbf{a}^*) \geq k(k - 1)$.

If there is a PSNE \mathbf{a}^* such that $SW_1(\mathbf{a}^*) \geq k(k - 1)$, the agents in 1 must form a clique of size k since the graph has a unit (positive) weight and the maximize social welfare of location 1 with capacity k is $k(k - 1)$. Thus, we have a clique of size k . \square

4 ALGORITHMS FOR COMPUTING A PSNE OF SG-LSI IN RESTRICTED SETTINGS

As discussed in the previous section, determining whether there is a PSNE in general is NP-complete. Our goal in this section is to develop efficient algorithms for computing a PSNE, if there exists one, in SG-LSI under some restricted settings.

4.1 Message-Passing Algorithms on Tree SG-LSI

We begin by developing a message-passing based algorithm for Tree SG-LSI when the underlying social network is a tree, regardless of the structure of the location graph. For the sake of simplicity, we first consider the case in which the capacity of each location is at least n .

More formally, we consider the following tree instance of SG-LSI

$$\mathcal{G}^T = \left(T, G^L, \{c_j\}_{j \in L}, S, \{\alpha_i\}_{i \in N}, \{\lambda_i\}_{i \in N}, \{b_{i,j}\}_{i \in N, j \in L}, f, \{u_i\}_{i \in N} \right),$$

where the social network among the agents is a tree $T = (N, E, w)$ rooted at $r \in N$. Although in our general model, f_i can be any general function of the agents in locations adjacent to i 's, for the purpose of applying dynamic programming here, we will assume that f_i depends only on those agents in adjacent locations that influence i . We assume $c_j = n$ for all $j \in L$. For any agent $i \neq r$, we use $pa(i)$ to denote i 's unique parent and $ch(i)$ to denote the set of i 's children.⁷ We next apply the Tree-Nash algorithm given by Kearns et al. [17] to our problem.

The message-passing algorithm has two phases: the upstream phase and the downstream phase. In the upstream phase, each non-root node i sends a message $T_{i \rightarrow j}(a_i, a_j) \in [0, 1]$ to its parent $j = pa(i)$, for each possible combination of location choices (a_i, a_j) . $T_{i \rightarrow j}(a_i, a_j) = 1$ if and only if there exists a *witness vector* $\{a_k\}_{k \in ch(i)}$ such that the following two conditions are satisfied.

(1) $T_{k \rightarrow i}(a_k, a_i) = 1$ for all $k \in ch(i)$, and

(2) a_i is i 's best response to parent j choosing a_j and each child k choosing a_k .

⁷We previously used $\mathcal{P}(i)$ to denote the set of i 's parents in the general social network.

The upstream phase begins with the leaf nodes and propagates upward to the root. Each leaf node l sends the following table to its parent $j = pa(l)$: $T_{l \rightarrow j}(a_l, a_j) = 1$ if and only if, for each $a'_l \in S$, $u_l(a_l, a_j) \geq u_l(a'_l, a_j)$. Note that the leaf nodes' messages are easy to compute due to the absence witness vectors.

Each internal node $i \neq r$ gathers messages $T_{k \rightarrow i}(a_k, a_i)$ from all $k \in ch(i)$ and constructs messages $T_{i \rightarrow j}(a_i, a_j)$ to send to the parent j . The main computational bottleneck comes from determining a witness vector $\{a_k\}_{k \in ch(i)}$. For general location effect function $f_i(N(a_i), \mathbf{a})$, we can go through all $m^{|ch(i)|}$ possible combinations of actions $\{a_k\}_{k \in ch(i)}$ to verify the two conditions for $T_{i \rightarrow j}(a_i, a_j) = 1$ stated two paragraphs ago. For any $T_{i \rightarrow j}(a_i, a_j) = 1$, it is sufficient for i to save just one witness vector corresponding to $T_{i \rightarrow j}(a_i, a_j)$ (all witness vectors need to be saved if we wish to compute all PSNE).

The downstream phase begins after the root node receives messages from its children. The root node r tries to find an action a_r for which there is a witness vector $\{a_k\}_{k \in ch(r)}$. If r cannot find any such action a_r , then there exists no PSNE. Otherwise, r chooses a_r and commands its children to choose actions according to the witness vector $\{a_k\}_{k \in ch(r)}$. Subsequently, each internal node i chooses the action a_i commanded by its parent and then commands its children to choose an action according to some witness vector for a_i . The process continues until all the leaf nodes have received messages from their parents. The choice of actions in the downstream phase constitutes a PSNE, if it exists. Like the original Tree-Nash algorithm, the running time of this algorithm is also exponential due to the exponential time spent on finding a witness vector at each internal node.

4.2 Efficient Message-Passing Algorithm on Tree SG-LSI with Additive Location Effect

We next present an efficient algorithm for finding a PSNE in a tree-structured SG-LSI when we have two locations and the location effect can be decomposed in an additive fashion as follows.

$$f_i(N(a_i), \mathbf{a}) = \sum_{j: (i,j) \in E \text{ and } (a_i, a_j) \in E^L} g_{i,j}(a_i, a_j).$$

Here, $g_{i,j}$ is a function that depends on the location choices a_i and a_j of agents i and j , respectively. We do not assume any particular functional form for $g_{i,j}$. The only assumption is that $g_{i,j}$ additively contributes to the location effect on i whenever j has an influence on i and j chooses a location in the neighborhood of i 's location.

We apply the same Tree-Nash framework here. However, adapting a technique used for 2-action tree influence games [16], we can now find a witness vector much more efficiently. For this, let us consider any internal node $i \neq r$ during the upstream phase. As usual, i gathers messages $T_{k \rightarrow i}(a_k, a_i)$ from all $k \in ch(i)$ and constructs messages $T_{i \rightarrow j}(a_i, a_j)$ to send to its parent j . Now, instead of going through all $m^{|ch(i)|}$ possible combinations of actions $\{a_k\}_{k \in ch(i)}$, we can find a witness vector more smartly. There are three cases.

Case I. There is some $k \in ch(i)$ such that $T_{k \rightarrow i}(a_k, a_i) = 0$ for all $a_k \in S$. In this case, there is no PSNE in the "subgame" downstream from i if i chooses a_i . As a result, i sends $T_{i \rightarrow j}(a_i, a_j) = 0$, for all $a_j \in S$.

Case II. For every $k \in ch(i)$, there is a unique action a_k such that $T_{k \rightarrow i}(a_k, a_i) = 1$. In this case there is just one candidate for a witness vector, which is $\{a_k\}_{k \in ch(i)}$. To make sure that it is in fact a witness vector, i verifies that a_i is a best response to parent j choosing a_j and each child k choosing a_k . In that case, i sends $T_{i \rightarrow j}(a_i, a_j) = 1$ and saves $\{a_k\}_{k \in ch(i)}$ as the witness vector. Otherwise, i sends $T_{i \rightarrow j}(a_i, a_j) = 0$.

Case III. The set of children of i , $ch(i)$ can be partitioned into two subsets: (1) S_u containing those children k that have a unique action a_k such that $T_{k \rightarrow i}(a_k, a_i) = 1$, and (2) $\overline{S_u}$ having those children k that have multiple values of a_k such that $T_{k \rightarrow i}(a_k, a_i) = 1$. For each $k \in \overline{S_u}$, among those a_k that make $T_{k \rightarrow i}(a_k, a_i) = 1$, we choose one particular a_k that maximizes the following expression.

$$\alpha_i w(k, i) \mathbb{1}[a_i = a_k] + \lambda_i g_{i,k}(a_i, a_k) \mathbb{1}[(a_i, a_k) \in E^L].$$

Equation 1 shows that the choosing a_k in this way bumps up i 's utility the most. Note that for any child $k \in S_u$, we do not have any choice other than to pick the unique a_k that makes $T_{k \rightarrow i}(a_k, a_i) = 1$. Next, we verify that a_i is i 's best response to parent j choosing a_j and each child k choosing a_k according to the above procedure. On successful verification, i sends $T_{i \rightarrow j}(a_i, a_j) = 1$ and saves the chosen $\{a_k\}_{k \in ch(i)}$ as the witness vector. On failure, we can be certain that there is no other choice of actions for any of i 's children that would lead to a greater utility for i , and as a result, i sends $T_{i \rightarrow j}(a_i, a_j) = 0$. The rest of the algorithm, including the downstream pass, is similar to the Tree-Nash algorithm described above. We obtain the following result.

THEOREM 4.1. *For any tree-structured SG-LSI with an additive location effect, two locations, and maximum indegree Δ , there exists an $O(n\Delta)$ algorithm for finding a PSNE or deciding there exists none.*

As a remark, to extend the above algorithm to more than two locations, it is not enough to maximize i 's utility at a_i in Case III above. We also need to simultaneously minimize i 's utilities at other locations, because depending on $\{a_k\}_{k \in \overline{S_u}}$, a_i may or may not become i 's best response. This does not happen in the two location case, because any $k \in \overline{S_u}$ can be either in a_i or the other location.

Algorithm 1: An Algorithm to Compute A PSNE

Input: SG-LSI with $p_1 \geq \dots \geq p_n$

Output: A PSNE profile \mathbf{a}^*

- 1 Let $\mathbf{a} = \mathbf{0}$ # Set \mathbf{a} to be a zero vector
 - 2 **for** $i = 1, \dots, n$ **do**
 - 3 Let $\mathcal{BR}_i(\mathbf{a}_{-i}) = \arg \max_{a_i \in S} u_i(a_i, \mathbf{a}_{-i})$.
 - 4 Select $j \in \mathcal{BR}_i(\mathbf{a}_{-i})$, set $a_i = j$.
 - 5 **end**
-

4.2.1 Special Case: Unlimited and Limited Capacity with no Location Effect. In this subsection, we study special cases of SG-LSI that can be solved in polynomial time. We first consider the instances of SG-LSI where $\lambda_i = 0$ for each $i \in N$, $b_{i,a_i} = 0$, and $w(i, i') = p_i p_{i'}$ for all $i, i' \in N$ and all locations a_i .

THEOREM 4.2. *There is a polynomial-time algorithm to find a PSNE of SG-LSI with $\lambda_i = 0$, $b_{i,a_i} = 0$, and $w(i, i') = p_i p_{i'}$ for all*

$i, i' \in N$ and all locations a_i . Here, $p_i > 0$ is some "personal value" of an agent i that helps factoring the influence weights.

PROOF. We first order the agents and location capacities such that $p_1 \geq \dots \geq p_n$ and $c_1 \geq \dots \geq c_m$, respectively. For $i = 1, \dots, m$ (in this order), assign the (remaining) highest c_i agents to location i until all of the agents have assigned some locations. It is not hard to see that the agents have no incentive to deviate. \square

We next consider the instances of SG-LSI where $\lambda_i = 0$ for each $i \in N$ and $w(i, i') = -p_i p_{i'}$ for all $i, i' \in N$.

THEOREM 4.3. *Algorithm 1 returns a PSNE of SG-LSI with $\lambda_i = 0$, $b_{i,a_i} = 0$, and $w(i, i') = -p_i p_{i'}$ for all $i, i' \in N$ and all locations a_i in polynomial time.*

5 EXPERIMENTAL RESULTS

In this section, our goal is to consider the impact of localized social influence and location effect on the PSNE in the SG-LSI. In our experiments, we count the total number of PSNE in randomly generated instances of SG-LSI. To account for the localized social influence and location effect, we vary different value of α_i and λ_i for each agent $i \in N$.

To generate an instance of SG-LSI, we use Erdos-Renyi (ER) random graphs as social networks. We first generate an ER graph by setting its parameter p (which specifies the probability of any two nodes will be connected by an edge). Once we have an ER graph, we randomly generate the weights of each edge to have a weight value between -1 and 1 (drawn uniformly at random). The value of $\alpha_i = \alpha$ and $\lambda_i = \lambda$ for each i are set such that $\alpha + \lambda_i = 1$. The utility of each agent is defined according to SG-LSI where the f_i is defined as in the Schelling Game [10]. One main difference with the Schelling Game is that "friends" of i are defined to be those with positive weighted edges to i while "enemies" are other agents (including the ones with negative weighed edges to i).

In the following, we consider $\alpha \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$, $n \in \{10, 20\}$, $m = 2$, and $p \in \{0.2, 0.5, 0.7\}$. For each combination of parameter values, we generate 10 random instances (of the edge weights). We report the number of PSNE for one of the instances below. We note that all other instances have a similar pattern in the results. The number of PSNE is computed using a brute-force approach to verify whether a strategy profile is a PSNE.

Figure 2 shows the number of PNSE for each combination of the parameters we considered. The x-axis represents the value of α , and the y-axis represents the number of PSNE. The left column corresponds to the case where $n = 10$, $p = 0.2$ (top left), $p = 0.5$ (middle left), and $p = 0.7$ (bottom left). The right column corresponds to the case where $n = 20$, $p = 0.2$ (top right), $p = 0.5$ (middle right), and $p = 0.7$ (bottom right). It is fairly easy to see that as we increase the α value (the term corresponds to localized social influence), the number of PSNE decreases for all combinations of parameters. In particular, for $\alpha = 0$ and $\lambda = 1$ (which imitates the Schelling Game [10]), we observe that there is a large number of PSNE compared to other combinations of $\alpha > 0$ and λ with localized social influence. Games with a large number of PSNE is often undesirable due to the lack of predictive capabilities. These experiments show that incorporating localized social influence provides a better predictive model by reducing the set of potential PSNE one needs to consider.

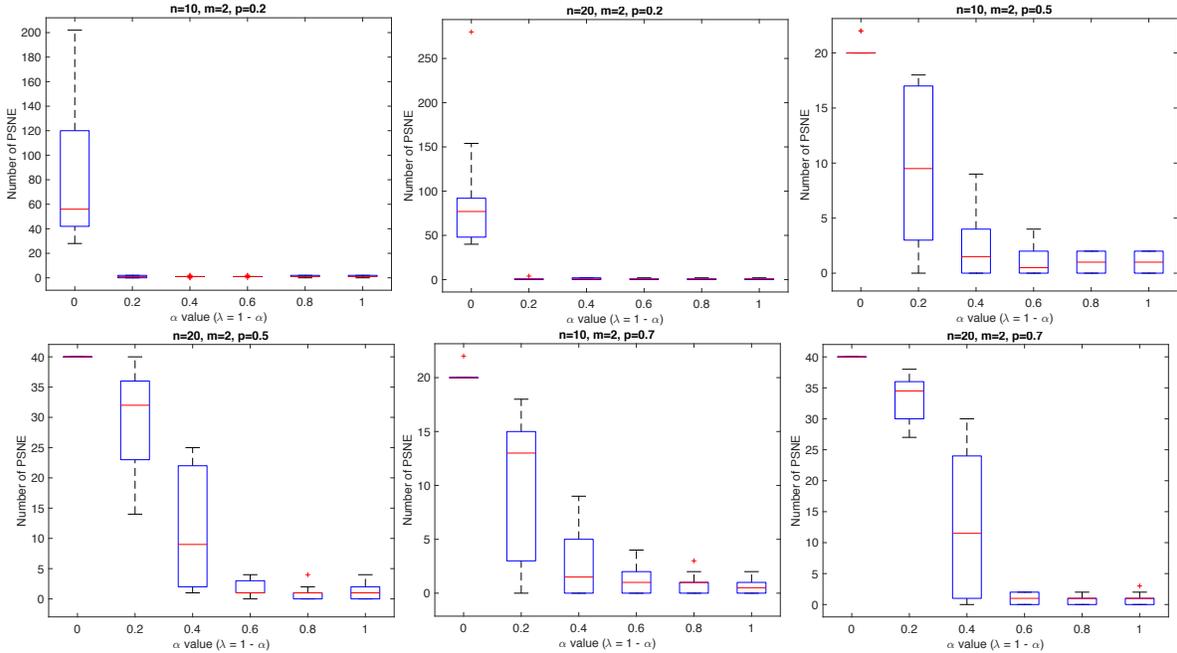


Figure 2: The Number of PSNE in Randomly Generated Instances of SG-LSI.

Top Left: $n=10, m=2, p=0.2$, Middle Left: $n=10, m=2, p=0.5$, Bottom Left: $n=10, m=2, p=0.7$

Top Right: $n=20, m=2, p=0.2$, Middle Right: $n=20, m=2, p=0.5$, Bottom Right: $n=20, m=2, p=0.7$

The x-axis represents the α value, and the y-axis represents the number of PSNE. Each plot corresponds to a combination of the parameters. The plots collectively show that, as we increase the value of α (i.e., the term corresponds to the localized social influence), the number of PSNE decreases.

6 CONCLUDING REMARKS

In this paper, we have introduced a game-theoretic framework for Schelling’s segregation where the agents’ choices of locations are intricately interdependent. Each agent’s choice depends on two factors: the location effect, which is inspired by the classical Schelling model and its extensions, and the localized social influence, which allows multiple agents to share a location while accounting for the complex interactions among them. We have shown that various special cases of our problems is provably hard. We have provided efficient algorithms for several special cases. Finally, our experimental results show the value of having the localized social influence component in Schelling models. There are several exciting future directions. First, we can use tree decomposition for general graphs and apply the adaptation of the Tree-Nash algorithm presented here. Formulating a message passing framework among the bags of vertices in tree decomposition is an interesting direction.

Another interesting direction is a thorough study of the efficiency of equilibria. We do have some preliminary results regarding the price of anarchy (PoA) and price of stability (PoS) [22], which we describe next. Very briefly, we can show that the PoA can be unbounded in general and the PoS can be 1 in a large class of SG-LSI.

In more details, if we consider only the location effect ($\alpha_i = 0$ and $\lambda_i = 1$ for all i) and take f_i to be the utility function of the standard Schelling games given in [10], then PoA can be shown to be unbounded as follows. Consider three agents a, b , and c , where

a and b influence each other positively (or are of the same type) and c is not connected to a or b (i.e., c is of a different type). Let the location graph be the line graph of three locations, each location having a capacity of 1. If we place c in the middle of the line graph, then the social welfare (which is the sum of the utilities of the agents) is zero. If we place either a or b in the middle of the line graph, then the social welfare is positive. Hence the PoA can be unbounded. Similar examples are used in [10].

If we consider only the localized social influence ($\alpha_i = 1$ and $\lambda_i = 0$ for all i), then PoA is also unbounded, as shown next. Consider again three agents a, b , and c , all of whom influence each other positively. Let the influence weights of the edges between a and b be $+\infty$ and the weights of the other edges be 1. Consider a location graph with only two connected locations, one with capacity 1 and the other with capacity 2. If we place a and b in the same location, then we obtain an optimal social welfare of $+\infty$. If we place a and b apart, then we obtain a social welfare of 2. Since either of these two scenarios is a PSNE, the PoA is unbounded, whereas the PoS is 1. In fact, the PoS is always 1 whenever the sum of the capacities is exactly equal to the number of agents for any instance of our game. We leave a thorough analysis as an interesting future direction.

ACKNOWLEDGEMENTS

We thank the anonymous referees for many helpful suggestions. MTI is grateful for support from NSF Award IIS-1910203.

REFERENCES

- [1] Aishwarya Agarwal, Edith Elkind, Jiarui Gan, and Alexandros A Voudouris. 2020. Swap Stability in Schelling Games on Graphs. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence*.
- [2] Christina Brandt, Nicole Immorlica, Gautam Kamath, and Robert Kleinberg. 2012. An analysis of one-dimensional Schelling segregation. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*. ACM, 789–804.
- [3] Ankit Chauhan, Pascal Lenzner, and Louise Molitor. 2018. Schelling segregation with strategic agents. In *International Symposium on Algorithmic Game Theory*. Springer, 137–149.
- [4] Wei Chen, Alex Collins, Rachel Cummings, Te Ke, Zhenming Liu, David Rincon, Xiaorui Sun, Yajun Wang, Wei Wei, and Yifei Yuan. 2011. Influence maximization in social networks when negative opinion may emerge and propagate. In *Proceedings of the Eleventh SIAM International Conference on Data Mining*. 379–390.
- [5] Nicholas A. Christakis and James H. Fowler. 2007. The Spread of Obesity in a Large Social Network over 32 Years. *N Engl J Med* 357, 4 (2007), 370–379. <https://doi.org/10.1056/NEJMs066082> arXiv:<http://content.nejm.org/cgi/reprint/357/4/370.pdf>
- [6] Nicholas A. Christakis and James H. Fowler. 2008. The Collective Dynamics of Smoking in a Large Social Network. *N Engl J Med* 358, 21 (2008), 2249–2258. <https://doi.org/10.1056/NEJMs0706154> arXiv:<http://content.nejm.org/cgi/reprint/358/21/2249.pdf>
- [7] William AV Clark. 1991. Residential preferences and neighborhood racial segregation: A test of the Schelling segregation model. *Demography* 28, 1 (1991), 1–19.
- [8] David A. Easley and Jon M. Kleinberg. 2010. *Networks, Crowds, and Markets - Reasoning About a Highly Connected World*. Cambridge University Press.
- [9] Hagen Echezell, Tobias Friedrich, Pascal Lenzner, Louise Molitor, Marcus Pappik, Friedrich Schöne, Fabian Sommer, and David Stangl. 2019. Convergence and Hardness of Strategic Schelling Segregation. In *Web and Internet Economics (WINE)*.
- [10] Edith Elkind, Jiarui Gan, Ayumi Igarashi, Warut Suksompong, and Alexandros A. Voudouris. 2019. Schelling Games on Graphs. (7 2019), 266–272. <https://doi.org/10.24963/ijcai.2019/38>
- [11] James H Fowler and Nicholas A Christakis. 2008. Dynamic spread of happiness in a large social network: Longitudinal analysis over 20 years in the Framingham Heart Study. *BMJ* 337, dec04_2 (2008), a2338–. <https://doi.org/10.1136/bmj.a2338> arXiv:http://www.bmj.com/cgi/reprint/337/dec04_2/a2338.pdf
- [12] Erez Hatna and Itzhak Benenson. 2012. The Schelling model of ethnic residential dynamics: Beyond the integrated-segregated dichotomy of patterns. *Journal of Artificial Societies and Social Simulation* 15, 1 (2012), 6.
- [13] Nicole Immorlica, Robert Kleinberg, Brendan Lucier, and Morteza Zadomighadam. 2017. Exponential segregation in a two-dimensional schelling model with tolerant individuals. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 984–993.
- [14] Mohammad T Irfan and Tucker Gordon. 2018. The Power of Context in Networks: Ideal Point Models with Social Interactions. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 910–918.
- [15] Mohammad T. Irfan and Luis E. Ortiz. 2011. A Game-Theoretic Approach to Influence in Networks. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*. 688–694. <http://www.aaai.org/ocs/index.php/AAAI/AAAI11/paper/view/3746>
- [16] Mohammad T. Irfan and Luis E. Ortiz. 2014. On influence, stable behavior, and the most influential individuals in networks: A game-theoretic approach. *Artificial Intelligence* 215 (2014), 79–119. <https://doi.org/10.1016/j.artint.2014.06.004>
- [17] M. Kearns, M. Littman, and S. Singh. 2001. Graphical Models for Game Theory. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*. 253–260.
- [18] David Kempe, Jon Kleinberg, and Éva Tardos. 2003. Maximizing the Spread of Influence through a Social Network. In *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- [19] Young Kim and Jaideep Srivastava. 2007. Impact of social influence in e-commerce decision making. In *Proceedings of the ninth international conference on Electronic commerce*. ACM, 293–302.
- [20] Alexander J Laurie and Narendra K Jaggi. 2003. Role of ‘vision’ in neighbourhood racial segregation: a variant of the Schelling segregation model. *Urban Studies* 40, 13 (2003), 2687–2704.
- [21] Jure Leskovec, Lada A Adamic, and Bernardo A Huberman. 2007. The dynamics of viral marketing. *ACM Transactions on the Web (TWEB)* 1, 1 (2007), 5.
- [22] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. 2007. *Algorithmic game theory*. Cambridge university press.
- [23] Thomas Schelling. 1978. *Micromotives and Macrobehavior*. Norton.
- [24] Thomas C Schelling. 1969. Models of segregation. *The American Economic Review* 59, 2 (1969), 488–493.
- [25] Thomas C. Schelling. 1971. Dynamic Models of Segregation. *Journal of Mathematical Sociology* 1, 2 (1971), 143 – 186.
- [26] H Peyton Young. 2009. Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning. *American economic review* 99, 5 (2009), 1899–1924.
- [27] Junfu Zhang. 2011. Tipping and residential segregation: a unified Schelling model. *Journal of Regional Science* 51, 1 (2011), 167–193.