

# Keeping Your Friends Close: Land Allocation with Friends

Edith Elkind<sup>1</sup>, Neel Patel<sup>2</sup>, Alan Tsang<sup>2</sup> and Yair Zick<sup>2</sup>

<sup>1</sup>University of Oxford

<sup>2</sup>National University of Singapore

elkind@cs.ox.ac.uk, {neel,atsang,zick}@comp.nus.edu.sg

## Abstract

We examine the problem of assigning plots of land to prospective buyers who prefer living next to their friends. In this setting, each agent’s utility depends on the plot she receives and the identities of the agents who receive the adjacent plots. We are interested in mechanisms without money that guarantee truthful reporting of *both* land values and friendships, as well as Pareto optimality and computational efficiency. We explore several modifications of the *Random Serial Dictatorship* (RSD) mechanism, and identify one that performs well according to these criteria. We also study the expected social welfare of the assignments produced by our mechanisms when agents’ values for the land plots are binary; it turns out that we can achieve good approximations to the optimal social welfare, but only if the agents value the friendships highly.

## 1 Introduction

Koranit, a village in the North Galilee region of Israel, was recently permitted to expand. Predetermined plots of land, of approximately equal size and price, have been drawn and must be assigned to prospective buyers<sup>1</sup>. However, while similar in size and official value, plots are not viewed as identical by the buyers: some buyers prefer living close to the village center, others favor living in an area with a view of the surrounding mountains, and yet others are interested in level plots suitable for a garden. Land ownership laws preclude direct ownership by buyers; rather, land is leased (for several decades) from a central governing body, and prospective buyers are prohibited by law from paying each other in order to secure land plots. In other words, land plots are to be treated as *indivisible goods*, allocated without monetary transfers. Prospective buyers form a small, close-knit community. Several of them are siblings (with parents having lived in the village for decades), or are long-term residents (in rented properties), with friends they’d like to be close to. Consequently, buyers have preferences not just over plots, but also over their potential neighbors. In fact, some pairs of buyers only care about being neighbors, regardless of where they

<sup>1</sup>One of these buyers happens to be the sister of the last author.




Figure 1: (a) Map of the proposed Koranit expansion (plots are numbered 1 – 35). Red lines denote roads. (b) The plot graph based on Figure 1a. Two plots are adjacent if they share a border.

end up. Thus, we are interested in *mechanisms that would enable the buyers to distribute the plots among themselves in a fair and efficient manner, and account for friendships.*

### 1.1 Our Contributions

We briefly discuss the complexity of finding an allocation that maximizes the social welfare in the complete information scenario, showing that this problem is NP-hard as well as hard to approximate. We then focus on the setting where each agent has at most one friend. This constraint is both realistic and simplifies our computational problem significantly: while our problem remains NP-hard, it admits a 2-approximation algorithm in this case.

We then investigate our problem from the perspective of mechanism design without money: can we incentivize the agents to truthfully report their plot values and friendship information? Given our application domain, we are interested in mechanisms that are simple to describe and participate in, while providing good social welfare guarantees. Since our problem generalizes the one-sided matching problem, a natural starting point is the Random Serial Dictatorship (RSD) mechanism, under which agents pick plots one by one. We show that RSD does not perform well in the presence of friendships, and explore several modifications of RSD, in which the picking order is based on friendship information. We identify settings in which our mechanisms are truthful and produce Pareto optimal outcomes, and provide bounds on their expected social welfare for the case where agents have binary valuations for the plots. For omitted proofs, see the full version of the paper [Elkind *et al.*, 2020].

## 1.2 Related Work

Well-known approaches to the one-sided matching problem include the Competitive Equilibrium from Equal Incomes (CEEI), proposed by Hylland and Zeckhauser [1979], the Probabilistic Serial (PS) mechanism [Bogomolnaia and Moulin, 2001], and Random Serial Dictatorship (RSD). However, neither CEEI nor PS are truthful; indeed, Svensson [1999] shows that RSD is the only truthful mechanism that satisfies ex-post Pareto optimality, anonymity and non-bossiness. Another intriguing truthful mechanism has been very recently proposed by Abebe *et al.* [2020].

The social welfare of truthful mechanisms in one-sided matching markets has been studied by Bhalgat *et al.* [2011] for rank-based valuation functions. Filos-Ratsikas *et al.* [2014] consider the social welfare of RSD for unit sum preferences, and show that RSD offers a  $\sqrt{n}$ -approximation to the optimal social welfare in this case. Adamczyk *et al.* [2014] focus on binary and unit-range preferences, and show that RSD offers a 3-approximation to the optimal social welfare for binary preferences, and a  $\sqrt{en}$ -approximation for unit range preferences. Christodoulou *et al.* [2016] analyze the Price of Anarchy (PoA) of one-sided matching markets for unit-sum preferences. They show that PoA for RSD is  $\mathcal{O}(\sqrt{n})$ . Krysta and Zhang [2016] study the one-sided matching market problem under matroid constraints. They propose a truthful mechanism and show that it offers  $\frac{e}{e-1}$ -approximation of the optimal social welfare.

Bodine-Baron *et al.* [2011] analyze a housing allocation problem where students have an inherent friendship structure. They focus on allocation stability and social welfare, rather than strategic behavior. An online variant of this problem is studied by Huzhang *et al.* [2017]. Massand and Simon [2019] also consider the stability of a one-sided matching market with externalities, but assume that agents cannot misreport their valuations.

## 2 Model and Preliminaries

We consider a set of agents  $N = \{1, \dots, n\}$  (land buyers) who need to be matched to  $n$  plots  $\mathcal{V} = \{v_1, \dots, v_n\}$ . Each agent receives exactly one plot. Thus, the goal is to output an *allocation*, i.e., a bijection  $A : N \rightarrow \mathcal{V}$ , where agent  $i$  gets plot  $A(i)$ .

We represent neighboring plots using a *plot graph*  $\mathcal{G} = \langle \mathcal{V}, \mathcal{E} \rangle$ : this is an undirected graph where *nearby* plots  $w$  and  $v$  are connected by an edge  $\{w, v\} \in \mathcal{E}$  (see Figure 1b). Each agent  $i \in N$  has a *valuation function*  $u_i : \mathcal{V} \rightarrow [0, 1] \cap \mathbb{Q}$ :  $u_i(v)$  is the value  $i$  derives from receiving plot  $v$ . Agents have friends, and care about living next to them. We represent friendships by a weighted directed *friendship graph*  $\langle N, F \rangle$ , where  $(i, j) \in F$  indicates that  $i$  and  $j$  are friends and the edge weight  $\varphi_{i,j} \in \mathbb{Q}_{\geq 0}$  is the additional utility  $i$  obtains for living next to  $j$ . We assume that friendships are reciprocal, but not necessarily symmetric; i.e.,  $(i, j) \in F \Leftrightarrow (j, i) \in F$ , but it may be the case that  $\varphi_{i,j} \neq \varphi_{j,i}$ . Let  $F^* = \{\{i, j\} : (i, j) \in F\}$ ; the unweighted undirected graph  $\langle N, F^* \rangle$  captures the presence of friendships, but not their weights. We set  $\varphi_{\min} = \min_{(i,j) \in F} \varphi_{i,j}$ . The quantity  $\varphi_{\min}$  plays an important role in our analysis: some of our

posed mechanisms offer better performance guarantees when  $\varphi_{\min} > 1$ , i.e. when the value of friendship exceeds the value of any plot.

The *utility*  $U_i(A)$  of agent  $i$  under allocation  $A$  is

$$u_i(A(i)) + \sum_{(i,j) \in F} \varphi_{i,j} \times \mathbb{I}(\{A(i), A(j)\} \in \mathcal{E}). \quad (1)$$

The first term in (1) is agent  $i$ 's utility from the plot she receives; the second term is her (non-negative) *externality* for nearby friends.

An instance of our allocation problem is a tuple  $I = \langle N, \mathcal{V}, \mathcal{E}, F, (u_i)_{i \in N}, (\varphi_{i,j})_{(i,j) \in F} \rangle$ ; let  $\mathcal{A}(I)$  denote the set of all allocations for an instance  $I$ .

The *social welfare* of an allocation  $A \in \mathcal{A}(I)$  is defined as the sum of agents' utilities:  $\text{SW}(A) = \sum_{i \in N} U_i(A)$ . Let  $\text{OPT}(I) = \max_{A \in \mathcal{A}(I)} \text{SW}(A)$ . Given two allocations  $A, A' \in \mathcal{A}(I)$ , we say that  $A'$  *dominates*  $A$  if  $U_i(A') \geq U_i(A)$  for all  $i \in N$ , and the inequality is strict for at least one agent. An allocation  $A$  is *Pareto optimal (PO)* if no other allocation dominates it. A non-PO allocation presents an avoidable loss of social welfare; we are thus interested in algorithms that output PO allocations.

We consider several special cases of our problem. We say that an instance  $I$  is *friendship-uniform* if there exists a positive value  $\varphi \in \mathbb{Q}_{\geq 0}$  such that  $\varphi_{i,j} = \varphi$  for all  $(i, j) \in F$ . We say that  $I$  is *binary* if  $u_i(v) \in \{0, 1\}$  for all  $i \in N, v \in \mathcal{V}$ . We say that  $I$  is *generic* if for every  $i \in N$ , every pair of plots  $v, w \in \mathcal{V}$  and every edge  $(i, j) \in F$  we have  $u_i(v) \neq u_i(w)$ ,  $u_i(v) \neq u_i(w) + \varphi_{i,j}$ . If each agent has at most one friend (an important assumption for the sequel), in a generic instance no agent is indifferent between two plots, even if one of them is adjacent to her friend's plot.

## 3 Optimal Friend-Constrained Allocations

We first analyze the complexity of finding (approximately) optimal allocations under the assumption of complete information, i.e., when the weighted friendship graph as well as agents' valuation functions are known. Formally, given an instance  $I$  of our problem and a positive rational value  $T$ , we ask whether there is an allocation  $A$  with  $\text{SW}(A) \geq T$ ; we refer to this problem as SW-OPT.

We first observe that even in the friendship-uniform case our problem is at least as hard as SUBGRAPH ISOMORPHISM, and hence NP-hard [Garey and Johnson, 1979]. To see this, let all agents value all plots at  $c \geq 0$  and each friendship at  $\varphi > 0$ ; the maximum social welfare achievable is  $n \times c + \varphi \times |F|$ . This welfare is obtained in allocations in which every pair of friends receive adjacent plots; such allocations exist if and only if  $\langle N, F^* \rangle$  is isomorphic to a subgraph of the plot graph  $\mathcal{G}$ . Thus, we obtain the following proposition.

**Proposition 3.1.** *SW-OPT is NP-complete. This result holds even if there exist  $c, \varphi > 0$  such that  $u_i(v) = c$  for all  $i \in N, v \in \mathcal{V}$  and  $\varphi_{i,j} = \varphi$  for all  $(i, j) \in F$ .*

The proof of Proposition 3.1 shows that SW-OPT is NP-hard even if the input instance is friendship-uniform and (i)  $\langle N, F^* \rangle$  consists of a single clique and a collection of singletons (in which case our problem is at least as hard as

CLIQUE), or (ii)  $\langle N, F^* \rangle$  has maximum degree 2 (in which case our problem is at least as hard as HAMILTONIAN CYCLE). The reduction from CLIQUE with  $c = 0$  also shows that SW-OPT is hard to approximate.

Motivated by these hardness/inapproximability results, in the remainder of the paper we focus on the setting where  $\langle N, F^* \rangle$  has maximum degree 1, i.e., it is a collection of edges (pairs of friends) and singleton nodes. In this case, the respective subgraph isomorphism problem reduces to finding a maximum matching in the plot graph, which can be done in polynomial time. While at a first glance this variant of the model may appear to be very restrictive, it is quite natural in our setting. Indeed, buying land is a serious commitment, so the ‘friendships’ in our context are typically sibling relationships, or other tight and long-running connections between households, and it is unlikely that a household is engaged in several such connections.

Nevertheless, even this special case of SW-OPT is NP-hard. The hardness result holds even in the friendship-uniform case and if the plot graph  $\mathcal{G}$  consists of a single path (i.e., plots are located along a road) and several isolated plots.

**Theorem 3.2.** *SW-OPT is NP-complete even if the instance is binary and friendship-uniform, the friendship graph  $\langle N, F^* \rangle$  has maximum degree 1, and the plot graph  $\mathcal{G}$  consists of a single path and isolated nodes.*

On the positive side, if the friendship graph has maximum degree 1, the problem of finding an allocation with maximum social welfare admits a poly-time 2-approximation algorithm.

**Theorem 3.3.** *Given an instance  $I$  where  $\langle N, F^* \rangle$  has maximum degree 1, we can compute in polynomial time an allocation  $A^*$  such that  $SW(A^*) \geq \frac{1}{2}OPT(I)$ .*

To summarize, for friendship graphs of maximum degree 1, SW-OPT is NP-hard, but admits a simple 2-approximation algorithm. In the remainder of the paper, we restrict ourselves to friendship graphs of maximum degree 1, and ask if this constraint allows us to find good allocations even when agents’ plot values and/or friendships are not publicly known.

## 4 Plot Allocation Mechanisms

In this section we adopt a *mechanism design* perspective. That is, we are interested in deterministic/randomized mechanisms (without money) that elicit valuations and friendships, and output an allocation based on the reports; these mechanisms should be simple to describe and participate in, and produce good allocations, even when agents are strategic.

We consider mechanisms where agents pick plots directly (and hence they do not need to report the plot values), and report their friendships; these reports may be used to select the picking order, and to restrict agents’ plot choices. Such mechanisms are easy to describe, making the allocation procedure transparent — an important concern in our setting.

Formally, we say that a deterministic mechanism is *friendship-truthful (FT)* if no agent can increase her utility by misreporting friendship information no matter what other agents report and no matter which plots they pick. A randomized mechanism is *universally FT* if it is friendship-truthful for every choice of its random bits, even when agents know

the random bits used by the mechanism. A deterministic mechanism is *Pareto optimal (PO)* if it is guaranteed to output a PO allocation on every input; a randomized mechanism is *universally PO* if it outputs a PO allocation on every input and for every choice of its random bits.

We are now ready to discuss mechanisms for land allocation with friends. We begin with serial dictatorship, identify its shortcomings, and explore several ways to overcome them. We derive a mechanism that is universally friendship-truthful, poly-time computable, and universally PO.

### 4.1 Serial Dictatorship

A natural starting point in our analysis is the (RANDOM) SERIAL DICTATORSHIP (RSD) mechanism [Abdulkadiroglu and Sonmez, 1998; Brandl *et al.*, 2016]. In the deterministic version of this mechanism, agents sequentially pick the plots, in a predetermined order; in the randomized version, agent order is chosen uniformly at random. For the one-sided matching problem, which is a special case of our problem, the optimal strategy of every agent under the SD mechanism is simple: she should simply choose the best available plot. Moreover, for one-sided matching the (R)SD mechanism is (universally) PO as long as all agents have generic utilities. Its performance with respect to the social welfare is well-understood; in particular, for binary utilities, a variant of this mechanism offers a constant-factor approximation to the optimal social welfare [Adamczyk *et al.*, 2014].

However, in the presence of friendships the agents’ decision problem under RSD becomes much more complicated, as illustrated by the following example.

**Example 4.1.** Consider an instance with agents 1, 2, 3, 4, and plots  $v_1, v_2, v_3, v_4$ , arranged on a path. Let  $F^* = \{\{1, 4\}, \{2, 3\}\}$ . Suppose that agents’ values for the plots are given by the table below and  $\varphi_{i,j} = .4$  for all  $(i, j) \in F$ . Consider what happens when we run the SD mechanism on this instance, with agent order  $(1, 2, 3, 4)$ .

	$v_1$	$v_2$	$v_3$	$v_4$
agent 1	.5	.3	0	0
agent 2	0	.5	.3	0
agent 3	0	.7	0	.5
agent 4	0	.5	0	0

Agent 1 picks first. If he were to choose  $v_1$ , agent 2 would face the choice between  $v_2$  and  $v_3$  ( $v_4$  is obviously less attractive). While she prefers  $v_2$ , she realizes that if she were to choose  $v_2$ , agent 3, who is her friend, would choose the non-adjacent plot  $v_4$ , so agent 2’s utility would be .5. If agent 2 chooses  $v_3$ , agent 3 would pick  $v_2$ , so agent 2’s utility would be  $.3 + \varphi_{2,3} = .3 + .4 = .7$ . Therefore, agent 2 picks  $v_3$ ; agent 3 picks  $v_2$  next, and finally agent 4 picks  $v_4$ . Under this scenario, agent 1 ends up several plots away from his friend, so his utility is .5.

Now, suppose that agent 1 chooses  $v_2$  instead. While  $u_1(v_2) = .3 < u_1(v_1)$ , in this case agent 2 would pick  $v_3$ , agent 3 would pick  $v_4$ , and agent 4 ends up with  $v_1$ , i.e., right next to agent 1. Thus, agent 1’s total utility from choosing  $v_2$  is .7. As his utility from choosing  $v_3$  or  $v_4$  is at most .4, his best choice is  $v_2$ , and the resulting allocation  $A$  is given by  $A(1) = v_2, A(2) = v_3, A(3) = v_4, A(4) = v_1$ .

Example 4.1 illustrates interesting phenomena that arise when using the SD mechanism. First, when deciding, agent 1 must reason about the decisions of agents who pick their plots after him. To choose optimally, he must know agents' plot values and friendships: indeed, if agents 3 and 4 had a low value for  $v_2$  and high values for  $v_3$  and  $v_4$ , he could safely pick  $v_1$ , as  $v_2$  would remain available for his friend.

Second, agent 1's decision depends on the order of agents who pick after him. If agent 4 picked immediately after agent 1, then agent 1 could safely pick  $v_1$  and expect agent 4 to pick the adjacent plot  $v_2$ . Consequently, his decision is even more difficult if the order of agents is unknown. In particular, if agent order is chosen uniformly at random (i.e., using RSD with hidden agent order), then, to evaluate the expected utility for each selection, he must consider all  $3! = 6$  scenarios corresponding to the permutations of the other agents.

Third, on this instance the SD mechanism produced an allocation that is not Pareto optimal: agents 1 and 2 would benefit from swapping their plots.

Thus, SD fails most of our criteria for a good mechanism. While it is simple to describe, the agents' decision problem is far from simple (in fact, the best upper bound on its computational complexity we could obtain is PSPACE). Further, agents must reason about other agents' utilities as well as their own, and the outcome may fail to be PO.

These difficulties mainly stem from the fact that whenever an agent  $i$  has a friend  $j$  that comes after her in the permutation,  $i$  must predict  $j$ 's decision. More specifically, for each available plot,  $i$  needs to know whether  $j$  can and will pick an adjacent plot on her turn. Clearly, this task is much easier when  $j$  picks immediately after  $i$ : indeed, in our example, agent 2 had a much easier time making up her mind than agent 1. Thus, we will now explore variants of the SD mechanism that enable friends to choose consecutively.

## 4.2 Choose-Together-SD (CT-SD) Mechanisms

Following the argument outlined at the end of Section 4.1, we consider a variant of RSD where if  $i$  and  $j$  are friends, they appear consecutively in the permutation.

ONLINE CHOOSE-TOGETHER RSD (ON-CT-RSD) is our first implementation of this idea. In this mechanism, at each step one of the unallocated agents is picked uniformly at random. The agent then picks a plot and may declare another unallocated agent as her friend; if she does, then her friend is the next to choose a plot (but cannot declare another friend). Let  $\mathcal{V}_i$  denote the set of plots that agent  $i$  can select from on her turn. We say that a plot  $v$  is a *singleton plot in  $\mathcal{V}_i$*  if it is not adjacent to any other plot in  $\mathcal{V}_i$ . If an agent  $i$  has a friend who has not selected a plot yet, and  $i$  selects a singleton plot in  $\mathcal{V}_i$ , then she will not be placed next to her friend in the resulting allocation.

Suppose first that the friendship information is publicly available, i.e., agent  $i$  can declare agent  $j$  to be her friend if and only if  $(i, j) \in F$ . In this case, under ON-CT-RSD the agents can compute their strategies in polynomial time.

**Theorem 4.2.** *Suppose that agents cannot misreport friendship information. Then each agent can compute her optimal strategy in polynomial time. To compute her strategy, each*

*agent only needs to know her preferences and the preferences of her friend (if she has one).*

A further appealing feature of ON-CT-RSD is that it is *ordinal*, in the sense that agents make their choice based on comparing plot values (accounting for additional value if a friend will be adjacent). However, even if friendships are public, ON-CT-RSD allocations are not necessarily PO; in fact, as Example 4.3 shows, ON-CT-RSD may output an allocation that is dominated by a better allocation.

**Example 4.3.** Consider an instance with agents 1, 2, 3 and plots  $v_1, v_2, v_3$ , where  $\mathcal{E} = \{\{v_2, v_3\}\}$ . Let  $F^* = \{\{1, 2\}\}$ . Agents' plot valuations are shown below, and  $\varphi_{1,2} = \varphi_{2,1} = .5$ . Suppose that the ON-CT-RSD mechanism picks agent 1 first, so the order is (1, 2, 3).

	$v_1$	$v_2$	$v_3$
agent 1	1	.9	0
agent 2	1	0	.4
agent 3	1	.1	0

Agent 1 can guarantee herself a utility of 1 by picking  $v_1$ . Her utility can be improved if she picked  $v_2$  and her friend, agent 2, cooperates by picking  $v_3$ . However, agent 2 would prefer  $v_1$  if it is available. Hence, the plot  $v_1$  remains agent 1's best choice, and the mechanism produces the allocation  $A(1) = v_1, A(2) = v_3, A(3) = v_2$ . Now, an allocation  $A'$  given by  $A'(1) = v_2, A'(2) = v_3, A'(3) = v_1$  dominates  $A$  with  $U_i(A') > U_i(A)$  for all  $i \in N$ .

Example 4.3 fails to produce a PO allocation: agent 2 does not choose a plot adjacent to her friend's because she gains more from choosing  $v_1$  over  $v_3$  than she gains from friendship. Indeed, if we change  $\varphi$  from .5 to 1, ON-CT-RSD produces a PO allocation. This observation can be generalized.

**Theorem 4.4.** *ON-CT-RSD is universally PO on generic instances  $\langle N, \mathcal{V}, \mathcal{E}, F, (u_i)_{i \in N}, (\varphi_{i,j})_{(i,j) \in F} \rangle$  with  $\varphi_{\min} > 1$ .*

So far we have assumed that agents cannot misreport their friendships. Let us now examine the role of this assumption.

**Proposition 4.5.** *ON-CT-RSD is not universally FT.*

*Proof.* Let us revisit Example 4.3. Suppose again that agent 1 is the first in the picking order. We argued that if agent 1 declares agent 2 as her friend, she maximizes her utility by picking the plot  $v_1$ , resulting in a total utility of 1. Suppose, however, that agent 1 picks plot  $v_2$  and declares agent 3 to be her friend. Then agent 3 chooses next, and picks the plot  $v_1$ . Agent 2 is then forced to pick plot  $v_3$ , so that the total utility of agent 1 is  $u_1(v_2) + \varphi_{1,2} = 1.4 > 1$ . Thus, agent 1 benefits from misreporting friendship information.  $\square$

However, as is the case for PO, if  $\varphi_{\min} > 1$ , this negative result no longer holds.

**Theorem 4.6.** *ON-CT-RSD is universally friendship-truthful for every instance with  $\varphi_{\min} > 1$ .*

To summarize, ON-CT-RSD is an attractive mechanism if  $\varphi_{\min} > 1$ ; however, in general it is neither universally PO nor universally friendship-truthful. We next discuss modifying this mechanism to avoid these issues.

### 4.3 Choose-Adjacent-SD (CA-SD) Mechanisms

The main reason why ON-CT-RSD fails both PO and friendship-truthfulness when  $\varphi_{\min} < 1$  is that when agent  $i$  declares agent  $j$  to be her friend,  $j$  can ‘jump the queue’, but may choose a plot not adjacent to  $i$ ’s. We now consider a mechanism that explicitly prohibits such behavior.

Specifically, this mechanism, ONLINE CHOOSE-ADJACENT RSD (ON-CA-RSD), proceeds identically to ON-CT-RSD with one difference: if agent  $i$  declares  $j$  to be her friend and chooses a non-singleton plot in  $\mathcal{V}_i$ , at the next step  $j$  *must* choose a plot adjacent to  $i$ ’s; if  $i$  chooses a singleton plot in  $\mathcal{V}_i$ ,  $j$  can then choose any plot in  $\mathcal{V}_j$ . Alternatively, if an agent selects a singleton plot, the mechanism may forbid her from declaring a friend; this has no impact on our analysis.

Note that ON-CA-RSD is equivalent to ON-CT-RSD if  $\varphi_{\min} > 1$ : whenever an agent  $i$  chooses after her friend, she would pick an adjacent plot if at all possible. However, in general, the mechanisms are different: e.g., on the instance described in Example 4.3 ON-CA-RSD would output an allocation  $A$  with  $A(1) = v_2, A(2) = v_3, A(3) = v_1$ .

It turns out that ON-CA-RSD satisfies the criteria formulated in the beginning of this section.

**Theorem 4.7.** *ON-CA-RSD is universally PO and universally friendship-truthful; moreover, agents’ strategies are polynomial-time computable.*

ON-CA-RSD has many attractive properties: it is simple, agents can compute their strategies efficiently and without knowing other agents’ preferences (not even their friends’ preferences!), and the mechanism always produces a PO allocation. However, if agents’ value for being close to their friends is low relative to the differences among the plot values, they may find this mechanism to be highly problematic.

**Example 4.8.** Let  $\mathcal{G}$  consist of a single edge  $\{v, w\}$  and  $n - 2$  isolated plots. Every agent values  $w$  at 0 and all other plots at 1. Suppose all friendships have value  $\varphi = .1$ . If agents  $i$  and  $j$  are friends and  $i$  is the first agent to pick, then  $i$  will choose  $v$  (as she can then benefit from being next to  $j$ ) and  $j$  will be forced to choose  $w$  and get the worst plot in  $\mathcal{V}$ .

One may then wonder if it is possible to modify ON-CA-RSD to give an agent the option to decline her friend’s ‘invitation’ and choose at a later point, but without having her plot choices constrained. There are several ways to implement this idea. For instance, if agent  $i$  declares a remaining agent  $j$  as a friend, we can offer  $j$  the choice of (1) picking a plot right after  $i$ , but it must be adjacent to  $i$ ’s plot (if at all possible), or (2) declining the invitation and returning to the pool of remaining agents; we refer to this mechanism as CA-BACK-TO-POOL-RSD (CA-BP-RSD). Alternatively, we can sample a default agent order in advance (uniformly among all possible  $n!$  orders), announce it to all agents, and then approach the agents one by one in this order, asking them to pick a plot and to declare a friend. If  $i$  declares  $j$  to be her friend, then  $j$  can either accept the invitation, jump the queue and pick a plot adjacent to  $i$ ’s (if such a plot exists); or, decline and keep her place in the queue (or, even more drastically,

move to the end of the queue); we refer to these mechanisms as CA-BACK-TO-QUEUE-RSD (CA-BQ-RSD) and CA-BACK-TO-END-RSD (CA-BE-RSD), respectively. These mechanisms seem to preserve the spirit of ON-CT-RSD, but offer agents more flexibility. Unfortunately, our next example shows that neither is universally friendship-truthful.

**Example 4.9.** Consider an instance with agents 1, 2, 3, 4, and plots  $v_1, v_2, v_3, v_4$ , arranged on a path in that order. Let  $F^* = \{\{1, 4\}\}$ . Suppose that agents’ values for the plots are given by the table below and  $\varphi_{1,4} = \varphi_{4,1} = .2$ .

	$v_1$	$v_2$	$v_3$	$v_4$
agent 1	0	1	0	0
agent 2	.3	0	.1	.2
agent 3	.3	0	.2	0
agent 4	0	0	0	1

Under ON-CA-RSD, if agent 1 picks first, she would pick  $v_2$ , and announce agent 4 as her friend, forcing agent 4 to pick an adjacent plot. Under CA-BP-RSD agent 4 can decline this option, in which case agents 2, 3, and 4 pick their plots in random order. Agent 4 chooses next w.p.  $1/3$ , in which case she will be able to pick her favorite plot. Thus, her expected utility is at least  $1/3 > \varphi_{4,1}$ , so she will not confirm friendship with agent 1. Thus, under CA-BP-RSD, if agent 1 declares agent 4 as her friend, her utility is 1.

Now, suppose agent 1 falsely declares agent 3 as her friend. Agent 3 has no reason to decline this invitation; indeed, accepting ensures that she receives her favorite plot (rather than risk losing it to agent 2). Thus, agent 3 accepts and picks  $v_1$ . Agents 2 and 4 prefer  $v_4$  to  $v_3$ , so the first to pick claims  $v_4$  for themselves. Thus, with probability .5 agent 4 ends up with  $v_3$ , which is adjacent to agent 1’s plot. Hence, under CA-BP-RSD, agent 1’s expected utility from declaring agent 3 as her friend is  $1 + .5 \times .2 = 1.1$ , which is higher than her utility from a truthful declaration.

The same argument shows that CA-BQ-RSD and CA-BE-RSD are not friendship-truthful: if the order is  $(1, 2, 3, 4)$ , then agent 1 prefers declaring agent 3 as her friend.

Thus, there does not seem to be an easy way to make ON-CA-RSD more flexible while retaining universal PO and friendship-truthfulness.

## 5 Social Welfare Maximization

So far, we focused on simplicity, polynomial-time computability and friendship-truthfulness; the only allocative efficiency measure we discussed was PO, which is a relatively weak requirement. We will now derive bounds on the social welfare of the assignments produced by ON-CT-RSD and ON-CA-RSD and their variants. For simplicity, we focus on friendship-uniform instances, i.e., we assume that  $\varphi_{i,j} = \varphi$  for some fixed  $\varphi$  and all  $(i, j) \in F$ . Since our problem is at least as hard as the one-sided matching problem, we cannot expect RSD and its variants to perform well for general valuations; thus, we focus on binary instances.

For binary utilities, Adamczyk *et al.* [2014] propose the following modification of the RSD mechanism, which we call RSD\*. In each iteration, before picking the next agent,

RSD\* asks all remaining agents to report if they have a positive value for some available plot. If some agents answer ‘yes’, RSD\* picks one of them uniformly at random, lets her pick a plot, and starts the next iteration; otherwise, RSD\* arbitrarily pairs remaining agents with remaining plots. RSD\* reduces waste while maintaining truthfulness, yielding a 1.45-approximation under binary valuations; can we obtain a similar approximation ratio in our setting?

Our first result is discouraging: ON-CT-RSD may produce assignments with very poor social welfare, even if  $\varphi > 1$ , i.e., even in the setting where it is PO for generic instances and friendship-truthful.

**Example 5.1.** Consider an instance with  $N = \{1, \dots, n\}$ ,  $\mathcal{V} = \{v_1, \dots, v_n\}$ , where  $\mathcal{E} = \{\{v_1, v_2\}\}$ . Suppose that  $F^* = \{\{1, 2\}\}$ ,  $\varphi_{1,2} = \varphi_{2,1} = 100$ . All agents value  $v_1$  at 1 and all other plots at 0.

Under ON-CT-RSD agents 1 and 2 end up in adjacent plots if and only if one of them appears first in the picking order, i.e., with probability  $\frac{2}{n}$ . Thus, the expected social welfare under this mechanism is  $1 + 2 \times \frac{2}{n} \times 100$ , whereas the optimal social welfare is 202.

When friendships are valuable, i.e.,  $\varphi \gg 1$ , we would like to avoid the situation described in Example 5.1. This can be accomplished by prioritizing pairs of friends, i.e., ensuring that pairs of friends choose first, followed by agents who do not have friends. This requires us to elicit friendship information offline, before agents start picking plots. As we cannot assume that agents will report this information truthfully, to fully specify such a mechanism, we need to handle inconsistent reports: what if  $i$  says that  $j$  is her friend, but  $j$  does not say that  $i$  is her friend? We take the conservative approach and treat  $i$  and  $j$  as friends iff both declare this friendship.

Formally, this mechanism, FRIENDS-FIRST CHOOSE-TOGETHER RSD\* (FF-CT-RSD\*) proceeds as follows. First, each agent reports who their friend is (or  $\emptyset$  for no friends). Let  $P$  be the set of pairs  $\{i, j\}$  who report each other as friends. We pick agents in the following order: as long as there exist a pair of adjacent unoccupied plots and  $P \neq \emptyset$ , we randomly remove a pair of agents  $\{i, j\}$  from  $P$ ;  $i$  and  $j$  then choose their plots (in random order). We execute RSD\* over remaining agents and plots once  $P = \emptyset$  or no adjacent plots are available. We analyze the performance of FF-CT-RSD\*, under the assumption that agents cannot lie about their friendships and  $\varphi > 1$ .

**Theorem 5.2.** *Let  $A$  be the output of FF-CT-RSD\* on a binary instance  $I$  with  $\varphi_{\min} > 1$ , where agents truthfully report friendships. Then  $\mathbb{E}(SW(A)) \geq \frac{1}{4}OPT(I)$ .*

Of course, since FF-CT-RSD\* prioritizes pairs of friends, we cannot expect it to be friendship-truthful. Thus, if friendship-truthfulness is considered desirable, we are left with ON-CA-RSD or its variants. Specifically, ON-CA-RSD, too, can be modified by pushing friendless agents who value all available plots at 0 to the back of the queue, in the spirit of RSD\*; we refer to this mechanism as ON-CA-RSD\*. It can be verified that this mechanism remains friendship-truthful.

Since ON-CA-RSD\* does not prioritize friendships, we cannot expect it to have a constant approximation ratio (con-

sider, e.g., its performance on the instance in Example 5.1). However, if  $\varphi > 1$ , we can bound the approximation ratio of ON-CA-RSD\* in terms of  $\varphi$ .

**Theorem 5.3.** *Let  $A$  be the output of ON-CA-RSD\* on a binary instance  $I$  with  $\varphi_{\min} > 1$ . Then  $\mathbb{E}(SW(A)) \geq \frac{1}{2\varphi+2}OPT(I)$ .*

The positive results presented so far in this section are for the case  $\varphi > 1$ . For  $\varphi < 1$ , positive results are more elusive. In particular, it is no longer the case that FF-CT-RSD\* has a constant approximation ratio.

**Proposition 5.4.** *There exists a friendship-uniform binary instance  $I$  with  $OPT(I) = 2 + 2\varphi$  such that the expected social welfare of the output of FF-CT-RSD is at most  $\frac{6}{n} + 4\varphi$ .*

Our last result applies not just to variants of the RSD mechanism, but to all truthful mechanisms: the approximation ratio of any such mechanism is at least  $1 + \frac{1}{2\varphi}$ , even if agents cannot misreport their friendship information.

**Proposition 5.5.** *Consider a mechanism  $\mathcal{M}$  that has access to the friendship graph  $\langle N, F \rangle$ , asks the agents to report their values for the plots, and outputs an allocation based on the agents’ report and the friendship graph. If no agent can benefit from misreporting her plot values under  $\mathcal{M}$  then here exists a friendship-uniform binary instance  $I = \langle N, \mathcal{V}, F, (u_i)_{i \in N}, (\varphi_{i,j})_{(i,j) \in F} \rangle$  such that for the allocation  $A$  output by  $\mathcal{M}$  we have  $\frac{\mathbb{E}(SW(A))}{OPT(I)} \leq \frac{2\varphi}{2\varphi+1}$ .*

## 6 Conclusions and Future Work

We analyze the problem of allocating plots to buyers who have intrinsic preferences over their neighbors. While the problem in its full generality offers several non-trivial computational challenges, we show that under some realistic assumptions on buyer preferences and permitted reports, it is possible to design simple mechanisms that maintain both truthful reporting and social welfare guarantees. We focus on RSD-like mechanisms for our problem; however, it may also be useful to consider other approaches, such as CEEI or PS, to check if they offer a tradeoff between truthfulness and social welfare. There are strong positive results if all agents value their friendships highly ( $\varphi_{i,j} > 1$ ), or not at all (i.e., if  $\varphi_{i,j} = 0$ ). Paradoxically, the presence of low-valued friendships may result in significant welfare loss, as shown by Proposition 5.5. This may be because even low-value friendships significantly distort agents’ behavior under RSD.

## Acknowledgements

This research was supported by an NRF Fellowship #R-252-000-750-733, an AI Singapore award #AISG-RP-2018-009, and the ERC Starting Grant ACCORD (GA 639945). Part of this work was done when the first author was visiting NUS. The authors thank the IJCAI 2020 and GAIW 2020 reviewers for their informative comments. Zick would like to thank his sister, Anat, for useful discussions and her insights.

## References

[Abdulkadiroglu and Sonmez, 1998] Atila Abdulkadiroglu and Tayfun Sonmez. Random serial dictatorship and the

- core from random endowments in house allocation problems. *Econometrica*, 66(3):689–702, May 1998.
- [Abebe *et al.*, 2020] Rediet Abebe, Richard Cole, Vasilis Gkatzelis, and Jason D. Hartline. A truthful cardinal mechanism for one-sided matching. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 2096–2113, 2020.
- [Adamczyk *et al.*, 2014] Marek Adamczyk, Piotr Sankowski, and Qiang Zhang. Efficiency of truthful and symmetric mechanisms in one-sided matching. In *International Symposium on Algorithmic Game Theory*, pages 13–24. Springer, 2014.
- [Bhalgat *et al.*, 2011] Anand Bhalgat, Deeparnab Chakrabarty, and Sanjeev Khanna. Social welfare in one-sided matching markets without money. In *Approximation, randomization, and combinatorial optimization. Algorithms and techniques*, pages 87–98. Springer, 2011.
- [Bodine-Baron *et al.*, 2011] Elizabeth Bodine-Baron, Christina Lee, Anthony Chong, Babak Hassibi, and Adam Wierman. Peer effects and stability in matching markets. In *International Symposium on Algorithmic Game Theory*, pages 117–129. Springer, 2011.
- [Bogomolnaia and Moulin, 2001] Anna Bogomolnaia and Hervé Moulin. A new solution to the random assignment problem. *Journal of Economic theory*, 100(2):295–328, 2001.
- [Brandl *et al.*, 2016] Florian Brandl, Felix Brandt, and Warut Suksompong. The impossibility of extending random dictatorship to weak preferences. *Economics Letters*, 141(C):44–47, 2016.
- [Christodoulou *et al.*, 2016] George Christodoulou, Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, Paul W Goldberg, Jie Zhang, and Jinshan Zhang. Social welfare in one-sided matching mechanisms. In *International Conference on Autonomous Agents and Multiagent Systems*, pages 30–50, 2016.
- [Elkind *et al.*, 2020] Edith Elkind, Neel Patel, Alan Tsang, and Yair Zick. Keeping your friends close: Land allocation with friends, 2020. ArXiv 2003.03558.
- [Filos-Ratsikas *et al.*, 2014] Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, and Jie Zhang. Social welfare in one-sided matchings: Random priority and beyond. In *International Symposium on Algorithmic Game Theory*, pages 1–12. Springer, 2014.
- [Garey and Johnson, 1979] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- [Huzhang *et al.*, 2017] Guangda Huzhang, Xin Huang, Shengyu Zhang, and Xiaohui Bei. Online roommate allocation problem. In *International Joint Conference on Artificial Intelligence*, pages 235–241, 2017.
- [Hylland and Zeckhauser, 1979] Aanund Hylland and Richard Zeckhauser. The efficient allocation of individuals to positions. *Journal of Political Economy*, 87(2):293–314, 1979.
- [Krysta and Zhang, 2016] Piotr Krysta and Jinshan Zhang. House markets with matroid and knapsack constraints. In *International Colloquium on Automata, Languages, and Programming*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2016.
- [Massand and Simon, 2019] Sagar Massand and Sunil Simon. Graphical one-sided markets. In *International Joint Conference on Artificial Intelligence*, pages 492–498, 2019.
- [Svensson, 1999] Lars-Gunnar Svensson. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, 16(4):557–567, 1999.