Keeping Your Friends Close: Land Allocation with Friends

Edith Elkind
University of Oxford
Oxford, United Kingdom
eelkind@gmail.com

Neel Patel
National University of Singapore
Singapore
neeltuk@gmail.com

Alan Tsang
National University of Singapore
Singapore
akhtsang@gmail.com

Yair Zick
National University of Singapore
Singapore
dcsyaz@nus.edu.sg

ABSTRACT

We examine the problem of assigning plots of land to prospective buyers who prefer living next to their friends. They care not only about the plot they receive, but also about their neighbors. This externality results in a highly non-trivial problem structure, as both friendship and land value play a role in determining agent behavior. We examine mechanisms that guarantee truthful reporting of *both* land values and friendships. We propose variants of *random serial dictatorship* (RSD) that can offer both truthfulness and welfare guarantees. Interestingly, our social welfare guarantees are parameterized by the value of friendship: if these values are low, enforcing truthful behavior results in poor welfare guarantees and imposes significant constraints on agents' choices; if they are high, we achieve good approximation to the optimal social welfare.

1 INTRODUCTION

A village in a quaint part of country X recently received a permission to expand. Predetermined plots of land, of approximately equal size and price, have been drawn and must be assigned to prospective buyers[1]. 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 amenable to a home garden. Land ownership laws preclude direct ownership by buyers; rather, land is leased 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*, and are to be 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 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.*

Figure 1: Map of the proposed village expansion (plots are numbered $1-35$). Red lines denote roads.

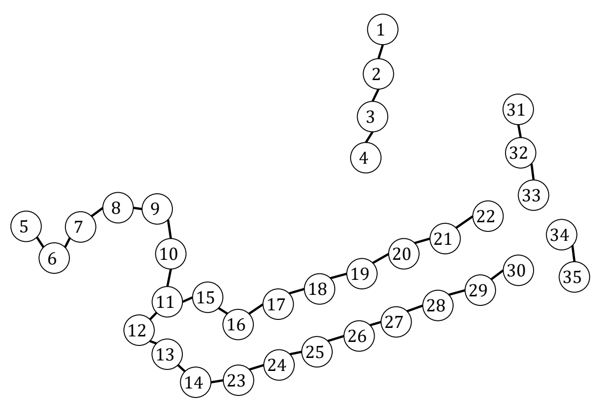

Figure 2: The plot graph based on Figure 1. Two plots are adjacent if they share a border

[1]As it happens, one of these buyers is the sister of an author.

Appears the 2nd Games, Agents, and Incentives Workshop (GAIW 2020). Held as part of the Workshops at the 19th International Conference on Autonomous Agents and Multiagent Systems., May 2020, Auckland, New Zealand., 2020

https://doi.org/doi

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. 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 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 [13], a natural starting point is the Random Serial Dictatorship (RSD) mechanism, under which agents pick plots one by one. We establish 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.

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 [13], the Probabilistic Serial (PS) mechanism [6], and Random Serial Dictatorship (RSD). However, neither CEEI nor PS are truthful; indeed, Svensson [17] 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. [2].

The social welfare of truthful mechanisms in one-sided matching markets has been studied by Bhalgat et al. [4] for rank-based valuation functions. Filos-Ratsikas et al. [10] 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. [3] 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. [9] analyze the Price of Anarchy (PoA) of one-sided matching markets for unit-sum preferences. They show that PoA for RSD is $O(\sqrt{n})$. Krysta and Zhang [14] 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. [5] 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. [12]. Massand and Simon [16] 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 omit several technical proofs from the paper due to page limits; these will appear in a full version of this work.

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 $A : N \to \mathcal{V}$, 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 1). Each agent $i \in N$ has a *valuation function* $u_i : \mathcal{V} \to [0, 1] \cap \mathbb{Q}$: $u_i(v)$ 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 proposed 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

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

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: $\mathrm{SW}(A) = \sum_{i\in N} U_i(A)$. Let $\mathrm{OPT}(I) = \max_{A\in\mathcal{A}(I)} \mathrm{SW}(A)$. 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 $i \in N$, $v \in V$. 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 $\mathrm{SW}(A) \geq T$; 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 [11]. 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 $\mathcal{G}$. This observation establishes 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 $i \in N$, $v \in V$ 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.*

Proof. It is immediate that this problem is in NP: we can compute the social welfare of a given allocation using formula (7). To prove hardness, we provide an NP-hardness reduction from Path Rainbow Matching.

An instance of the Path Rainbow Matching problem is given by an integer k, and a properly edge-colored path, i.e., an undirected path $P = \langle V, E \rangle$ with vertices $V = \{v_1, \dots, v_s\}$, and edges $E = \{e_1, \dots, e_{s-1}\}$ such that $e_i = \{v_i, v_{i+1}\}$ for all $i = 1, \dots, s-1$ together with a finite set of colors $C = \{c_1, \dots, c_q\}$, and a mapping $\xi : E \to C$ from edges of P to colors such that $\xi(e_i) \neq \xi(e_{i+1})$ for each $i = 1, \dots, s-1$. An instance $\langle k, P, C, \xi \rangle$ is a 'yes'-instance if there exists a subset of edges $M \subseteq E$ with $|M| \geq k$ such that all edges in M are pairwise disjoint and have different colors, and a 'no'-instance otherwise. This problem is known to be NP-hard [15].

Given an instance of Path Rainbow Matching $\langle k, P, C, \xi \rangle$, we construct an instance $\langle N, \mathcal{V}, \mathcal{E}, F, (u_i)_{i \in N}, (\varphi_{i,j})_{(i,j) \in F} \rangle$ of our problem as follows.

We set $N = \{1, \dots, 2q + s\}$: there are two agents corresponding to each color and s additional dummy agents. We refer to agents $1, \dots, 2q$ as *color agents*.

The plot graph $(\mathcal{V}, \mathcal{E})$ has $\mathcal{V} = \{v_1, \dots, v_s, w_1, \dots, w_{2q}\}$, and $\mathcal{E} = \{\{v_i, v_{i+1}\} : i = 1, \dots, s-1\}$, i.e., it is a copy of the given path instance together with $2q$ additional isolated plots (one for each color agent).

The agents are friends if and only if they correspond to the same color and all friendships have weight $\varphi = .1$: we set $F = \{(2i-1, 2i), (2i, 2i-1) : 1 \leq i \leq q\}$ and $\varphi_{i,j} = .1$ for each $(i, j) \in F$.

The agents' plot valuations are defined as follows. Dummy agents value all plots at 0: $u_i(x) = 0$ for each $i = 2q+1, \dots, 2q+s$ and each $x \in \mathcal{V}$. Each color agent values 'her' isolated plot at 1 and all other isolated plots at 0: for each $i = 1, \dots, 2q$ we have $u_i(w_i) = 1$, $u_i(w_j) = 0$ for $j \neq i$. Also, for each edge $\{x, y\}$ with $\xi(\{x, y\}) = c_i$ one agent in $\{2i-1, 2i\}$ values x at 0 and y at 1, and the other agent values y at 0 and x at 1. Specifically, for every color $c_i \in C$ let $\mathcal{E}_i = \{e \in \mathcal{E} : \xi(e) = c_i\}$, and suppose that $\mathcal{E}_i = \{\{v_{i_1}, v_{j_1}\}, \dots, \{v_{i_r}, v_{j_r}\}\}$, where $i_1 < i_2 < \cdots < i_r$ and $j_\ell = i_\ell + 1$ for $\ell = 1, \dots, r$. Then agent $2i-1$ values a plot $v_k \in \{v_1, \dots, v_s\}$ at 1 if $k = i_\ell$ for an odd value of ℓ or $k = j_\ell$ for an even value of ℓ, and otherwise she values it at 0. Similarly, agent $2i$ values a plot $v_k \in \{v_1, \dots, v_s\}$ at 1 if if $v = i_\ell$ for an even value of ℓ or $v = j_\ell$ for an even value of j, and otherwise she values it at 0.

We claim that our instance admits an assignment A with $\mathrm{SW}(A) \geq 2q + 2k\varphi$ if and only if $\langle k, P, C, \xi \rangle$ is a 'yes'-instance of Path Rainbow Matching. Indeed, suppose we start with a 'yes'-instance of Path Rainbow Matching, and let M be the respective matching. For each edge $e = \{u, v\} \in M$, if $\xi(e) = c_i$, we assign color agents $2i-1$ and $2i$ to the endpoints of e so that each of them is given the endpoint that she values at 1; we assign the remaining color agents to their preferred isolated plots, while the dummy agents are matched arbitrarily to the remaining nodes. Then each color agent values her plot at 1, and in addition there are k pairs of friends who are allocated adjacent plots, so the overall social welfare is $2q + 2k\varphi$.

Conversely, suppose that there is an allocation A with $\mathrm{SW}(A) \geq 2q + 2k\varphi$. Suppose first there exists some color agent i with $U_i(A) < 1$. Then $A(j) = w_i$ for some agent $j \in N \setminus \{i\}$, and $u_j(w_i) = 0$, so we can swap i and j and increase the overall social welfare: even if $A(i)$ is adjacent to the plot of i's friend, the loss in social welfare caused by moving i away from her friend is at most $2\varphi < 1$, and the gain in plot values is 1. Thus, we can assume that in A each color agent values her plot at 1. This means that there exist at least k pairs of friends who are allocated adjacent plots, with each friend valuing her plot at 1. Let $2i-1$, $2i$ be some such pair of friends, and suppose that they have been allocated plots v_j, v_{j+1}. Then either $\xi(\{v_j, v_{j+1}\}) = c_i$ or $\xi(\{v_{j-1}, v_j\}) = \xi(\{v_{j+1}, v_{j+2}\}) = c_i$. However, the latter case is impossible: we defined the plot valuation functions so that at least one of the agents $2i-1$ and $2i$ values both v_j and v_{j+1} at 0 in this case. Thus, these pairs correspond to a rainbow matching in P of size at least k. □

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 $SW(A^*) \geq \frac{1}{2}OPT(I)$.*

Proof. Our algorithm proceeds as follows.

First, we need to find a maximum matching in the graph $\langle \mathcal{V}, \mathcal{E}\rangle$; let $\{\{v_1, w_1\}, \ldots, \{v_s, w_s\}\}$ be the set of edges of this matching. Suppose that $F^* = \{\{i_1, j_1\}, \ldots, \{i_t, j_t\}\}$, where $\varphi_{i_\ell, j_\ell} + \varphi_{j_\ell, i_\ell} \geq \varphi_{i_r, j_r} + \varphi_{j_r, i_r}$ whenever $1 \leq \ell < r \leq t$, i.e., the edges in F^* are sorted by the total friendship weight in non-increasing order. Then for each $k = 1, \ldots, \min\{t, s\}$ we allocate plot v_k to i_k and plot w_k to j_k; all remaining plots are matched arbitrarily to the remaining agents. Let the resulting allocation be A_1.

Second, we consider the weighted complete bipartite graph with parts N and $\mathcal{V}$ where the weight of an edge $\{i, v\} \in N \times \mathcal{V}$ is equal to $u_i(v)$, and compute an allocation that corresponds to a maximum-weight matching in this graph; let this allocation be A_2.

We output the better of the two allocations A_1 and A_2 (breaking ties arbitrarily). To see that this algorithm provides a $\frac{1}{2}$-approximation, consider an arbitrary allocation A. Under this allocation, at most $\min\{s, t\}$ pairs of friends are allocated adjacent plots, so the total utility they derive from friendship is at most $\mathrm{SW}(A_1)$. Moreover, the total value that the agents assign to their plots under A is at most $\mathrm{SW}(A_2)$. Thus, we have $\mathrm{SW}(A) \leq \mathrm{SW}(A_1) + \mathrm{SW}(A_2)$, whereas we output an allocation whose social welfare is at least $\frac{1}{2}(\mathrm{SW}(A_1) + \mathrm{SW}(A_2))$. Moreover, both A_1 and A_2 can be computed in polynomial time, which concludes the proof. □

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: the simplest such mechanism is the Serial Dictatorship mechanism, where agents sequentially pick plots. Given that agents' utilities may depend on what other agents pick in subsequent rounds, we require that agents should be able to efficiently compute optimal strategies. In addition, we require agents to report friendships. These reports are used to select the picking order, and to possibly 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 [1, 7]. 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 [3].

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 2'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 2 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's turn follows 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: at each step, the mechanism picks one of the unallocated agents uniformly at random. The agent then picks a plot and may declare another unallocated agent as her friend; if so, 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).*

Proof. Since agents cannot misreport friendships, their strategic decisions are limited to what plot to pick. Consider an agent $i \in N$. If i has no friends, she should simply pick the plot with the highest value among the available plots. Now, suppose that $(i, j) \in F$. If i picks after j, she can choose the plot that maximizes her utility, taking her friend's (known) location into account. Finally, if i picks before j, she can consider all available plots, and, for each option, check whether j would choose one of the adjacent plots at the next step; in this way, she can determine which plot would maximize her utility. Note that to make her decision, i does not have to reason about the utilities of agents in $N \setminus \{i, j\}$. □

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. On-CT-RSD may produce an allocation A where another allocation A' exists with $U_i(A') > U_i(A)$ for all $i \in N$, even if agents cannot misreport friendships.

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 φ 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* $\varphi_{\min} > 1$.

Proof. Suppose for the sake of contradiction that, given an instance of our problem with $\varphi_{\min} > 1$, On-CT-RSD produces an allocation A, yet there exist another allocation A' for this instance such that $U_\ell(A') \geq U_\ell(A)$ for all $\ell \in N$ and $U_i(A') > U_i(A)$ for some $i \in N$. We can assume without loss of generality that under On-CT-RSD the picking order is $(1, 2, \ldots, n)$. Let i be the first agent in this order such that $U_i(A') > U_i(A)$; note that, since our instance is generic, this means that $A(\ell) = A'(\ell)$ for all $\ell < i$ and hence $A'(i) \in \mathcal{V}_i$.

Suppose first that i has no friends. Then under On-CT-RSD she picks the most valuable plot in $\mathcal{V}_i$ and $A'(i) \in \mathcal{V}_i$, so we have $U_i(A) = u_i(A(i)) \geq u_i(A'(i)) = U_i(A')$, a contradiction. Thus, we can assume that $(i, j) \in F$ for some $j \in N$.

Now, suppose that in our run of On-CT-RSD agent i picks her plot after j, and hence $A'(j) = A(j)$. Then the utility that i would obtain by picking $A'(i)$ in the execution of On-CT-RSD is equal to the utility she obtains in A'; since $A'(i) \in \mathcal{V}_i$, we obtain a contradiction again.

It remains to consider the case where in our run of On-CT-RSD agent i chooses before agent j (and then j chooses next). Then i's best strategy is to pick the highest-value non-singleton plot in $\mathcal{V}_i$ (and to simply pick the highest-value plot if all plots in $\mathcal{V}_i$ are singletons). Indeed, if i picks a non-singleton plot in $\mathcal{V}_i$, since $\varphi_{j,i} > 1$, agent j would necessarily pick an adjacent plot in the next step, and, since $\varphi_{i,j} > 1$, agent i would derive a higher utility from this choice than from any singleton plot in $\mathcal{V}_i$.

Suppose first that all plots in $\mathcal{V}_i$ are singletons, and hence under On-CT-RSD agent i picks the highest-value singleton plot in $\mathcal{V}_i$. Then $A'(i)$, too, is a singleton plot in $\mathcal{V}_i$, i.e., in A all plots adjacent to $A'(i)$ are occupied by agents who appear before i in the picking order, and we know that these agents are allocated the same plots in A'. Thus, i and j are not allocated adjacent plots in A', and hence $U_i(A') = u_i(A'(i)) \leq u_i(A(i)) = U_i(A)$, where the inequality holds since under On-CT-RSD agent i picks the highest-value singleton plot in $\mathcal{V}_i$. Thus, we obtain a contradiction in this case.

To complete the proof, suppose that i picks a non-singleton plot in $\mathcal{V}_i$ under On-CT-RSD, and therefore $U_i(A) > 1$. Then it has to be the case that $U_i(A') > 1$, i.e., in A' agents i and j are allocated adjacent plots. Thus, $A'(i)$ is a non-singleton plot in $\mathcal{V}_i$, but then we obtain a contradiction again, since $A(i)$ is the highest-value non-singleton plot in $\mathcal{V}_i$, so $U_i(A) = u_i(A(i)) + \varphi_{i,j} \geq u_i(A'(i)) + \varphi_{i,j} = U_i(A')$. □

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. □

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* $\varphi_{\min} > 1$.

Proof. Clearly, if an agent has no friends, she cannot benefit from declaring another agent to be her friend, as it would not give her access to a better plot. Similarly, if an agent j is 'invited' by i, i.e., j picks right after i because i declared j to be her friend, j is not asked to report her friendship information, so she has no opportunity to misreport. Now, suppose that i has a friend (say, j), and i gets to pick a plot before j. If all plots in $\mathcal{V}_i$ are singletons, then friendship information is irrelevant, and i has no incentive to misreport. Otherwise, let v be a highest-value non-singleton plot in $\mathcal{V}_i$. Then the highest utility i can hope to get in this run of the mechanism is $u_i(v) + \varphi_{i,j}$, which is exactly the utility she would get by picking v and declaring j to be her friend: indeed, since $\varphi_{j,i} > 1$, j would then choose a plot adjacent to v. Hence, i has no incentive to misreport the friendship information in this case as well. □

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.*

Proof. The analysis is similar to the analysis for On-CT-RSD with $\varphi > 1$. Suppose for the sake of contradiction that, given an instance of our problem, On-CA-RSD produces an allocation A, yet there exist another allocation A' for this instance such that $U_\ell(A') \geq U_\ell(A)$ for all $\ell \in N$ and $U_i(A') > U_i(A)$ for some $i \in N$. We can assume without loss of generality that under On-CA-RSD the picking order is $(1, 2, \ldots, n)$. Let i be the first agent in this order such that $U_i(A') > U_i(A)$; since or instance is generic, this means that $A(\ell) = A'(\ell)$ for all $\ell < i$ and hence $A'(i) \in \mathcal{V}_i$.

If i has no friends, then under On-CA-RSD she picks the most valuable plot in $\mathcal{V}_i$; as $A'(i) \in \mathcal{V}_i$, we have $U_i(A) = u_i(A(i)) \geq u_i(A'(i)) = U_i(A')$, a contradiction.

Now, suppose that $(i, j) \in F$ for some $j \in N$. If in our run of On-CA-RSD agent i picks her plot after j, we have $A'(j) = A(j)$.

Suppose first that agent j picked a singleton plot in $\mathcal{V}_j$, so that the choice of agent i is unconstrained, and hence she picks the most valuable plot in $\mathcal{V}_i$. Then the analysis is similar to the previous case: as $A'(i) \in \mathcal{V}_i$, we have $U_i(A) = u_i(A(i)) \geq u_i(A'(i)) = U_i(A')$, a contradiction. On the other hand, if j did not pick a singleton plot, then $A(i)$ is the most valuable plot among the plots that are adjacent to $A(j)$. Thus, if $A'(i) \neq A(i)$ then $A'(i)$ is not adjacent to $A'(j) = A(j)$. But this means that $U_j(A') = U_j(A) - \varphi_{j,i}$, a contradiction with our assumption that $U_\ell(A') \geq U_\ell(A)$ for all $\ell \in N$.

Finally, suppose that in our run of On-CA-RSD agent i picks her plot before j. Note that $A'(i), A'(j) \in V_i$. If these plots are adjacent, then i can pick $A'(i)$; as j will be forced to pick an adjacent plot in the next iteration, we have $U_i(A) \geq U_i(A')$. Otherwise, we have $U_i(A') = u_i(A'(i))$, so agent i can obtain the same utility as in A' simply by choosing $A'(i)$. This completes the proof of Pareto optimality.

For friendship truthfulness, the proof is very similar to the proof of Theorem 4.6: just as in that proof, an agent does not benefit from misreporting if she does not have a friend or if she chooses after her friend. Further, if agent i chooses before her friend j, the maximum utility she can obtain is the higher of $\max_{v\in\mathcal{V}_i} u_i(v)$ and $\max_{v\in\mathcal{V}_i^{ns}} u_i(v) + \varphi_{i,j}$, where $\mathcal{V}_i^{ns}$ is the set of non-singleton plots in $\mathcal{V}_i$, and she can guarantee herself that utility by reporting truthfully.

Finally, the polynomial-time computability follows from the description of the optimal strategies given in the previous paragraph. □

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 φ 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. [3] 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 and terminates. RSD* reduces waste while maintaining truthfulness, giving a 1.45-approximation to the optimal social welfare 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 then 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 ∅ 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 (consider, 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 φ.

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 $\frac{6}{n} + 4\varphi$.*

Proof. Consider an instance with $N = \{1, \dots, n\}$, where $n = 2k$ is even, $\mathcal{V} = \{v_1, \dots, v_k, w_1, \dots, w_k\}$, $\mathcal{E} = \{\{v_1, v_i\} : 2 \leq i \leq k\} \cup \{\{w_1, w_i\} : 2 \leq i \leq k\} \cup \{\{v_1, w_1\}\}$, $F^* = \{\{2i-1, 2i\} : i = 1, \dots, k\}$. Suppose that $u_1(v_1) = u_1(w_1) = u_2(v_1) = u_2(w_2) = 1$ and all other plot values are 0.

If $\varphi < .5$, an optimal allocation assigns v_1 and w_1 to agents 1 and 2, so that the social welfare is $2 + 2\varphi$. Now, under FF-CT-RSD* the probability that agents 1 and 2 appear in the first two positions of the picking order is $\frac{2}{n}$, and the probability that they appear in the next two positions of the picking order is $\frac{2}{n}$ as well; if neither of these events happens, plots v_1 and w_1 will be occupied by agents who value them at 0 (but derive positive utility from being next to their friend), so the social welfare will be at most 4φ. Thus, the expected social welfare of the allocation produced by FF-CT-RSD* is at most $2 \times \frac{2}{n} + 1 \times \frac{2}{n} + 4\varphi$. □

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 most $1 + \frac{1}{2\varphi}$, 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 $\frac{\mathbb{E}(SW(A))}{OPT(I)} \leq \frac{2\varphi}{2\varphi+1}$.*

Proof. Let $I_1 = \langle N, \mathcal{V}, F, (u_i)_{i \in N}, (\varphi_{i,j})_{(i,j) \in F} \rangle$, where $N = \{1, \dots, n\}$ and n is even, $n = 2k$, $\mathcal{V} = \{v_1, \dots, v_n\}$, $\mathcal{E} = \{\{v_1, v_i\} : 2 \leq i \leq n\}$, $F = \{\{2i-1, 2i\} : 1 \leq i \leq k\}$, $u_i(v) = 0$ for all $i \in N$ and all $v \in \mathcal{V}$, and there exists a positive value φ such that $\varphi_{i,j} = \varphi$ for all $(i, j) \in F$. That is, the plot graph is a star with center v_1, and each agent has a friend and values all plots at 0. We have $\mathrm{SW}(A) = 2\varphi$ for every $A \in \mathcal{A}(I_1)$.

By the pigeonhole principle, there exists a pair of friends $\{2i - 1, 2i\}$ such that mechanism $\mathcal{M}$ allocates v_1 to $2i - 1$ or $2i$ with probability at most $\frac{2}{n}$. Now, consider the instance I_2 that is obtained from I_1 by changing $u_{2i}(v_1)$ to 1. Since $\mathcal{M}$ is truthful, agent $2i$ cannot increase her utility in I_1 by misreporting her utility function, so given I_2, $\mathcal{M}$ allocates v_1 to $2i$ with probability at most $\frac{2}{n}$. Thus, the expected social welfare of the allocation produced by $\mathcal{M}$ on I_2 is at most $\frac{2}{n} + 2\varphi$, whereas $\mathrm{OPT}(I_2) = 1 + 2\varphi$. As n can be arbitrarily large, the bound follows. □

6 CONCLUSIONS AND FUTURE WORK

We have analyzed the problem of allocating plots of land 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 obtain positive results if all agents value their friendships highly ($\varphi_{i,j} > 1$), and even stronger positive results are known in the absence of friendships (i.e., if $\varphi_{i,j} = 0$). However, paradoxically, the presence of low-valued friendships may result in significant welfare loss, as shown by Proposition 5.5. To see why this may be the case, note that even low-value friendships may distort agents' behavior under RSD, thereby changing the allocation significantly.

We focused on RSD-like mechanisms for our problem; however, it may also be useful to consider other approaches. E.g., we can explore market-like mechanisms, where agents are allocated identical budgets and need to bid on plots and possibly on friendships, in the spirit of Budish [8].

ACKNOWLEDGEMENTS

This research was supported by an NRF Fellowship #R-252-000-750-733, an AI Singapore award #AISG-RP-2018-009, and th 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

[1] Atila Abdulkadiroglu and Tayfun Sonmez. 1998. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* 66, 3 (May 1998), 689–702.

[2] Rediet Abebe, Richard Cole, Vasilis Gkatzelis, and Jason D. Hartline. 2020. A Truthful Cardinal Mechanism for One-Sided Matching. In *ACM-SIAM Symposium on Discrete Algorithms*. 2096–2113.

[3] Marek Adamczyk, Piotr Sankowski, and Qiang Zhang. 2014. Efficiency of truthful and symmetric mechanisms in one-sided matching. In *International Symposium on Algorithmic Game Theory*. Springer, 13–24.

[4] Anand Bhalgat, Deeparnab Chakrabarty, and Sanjeev Khanna. 2011. Social welfare in one-sided matching markets without money. In *Approximation, randomization, and combinatorial optimization. Algorithms and techniques*. Springer, 87–98.

[5] Elizabeth Bodine-Baron, Christina Lee, Anthony Chong, Babak Hassibi, and Adam Wierman. 2011. Peer effects and stability in matching markets. In *International Symposium on Algorithmic Game Theory*. Springer, 117–129.

[6] Anna Bogomolnaia and Hervé Moulin. 2001. A new solution to the random assignment problem. *Journal of Economic theory* 100, 2 (2001), 295–328.

[7] Florian Brandl, Felix Brandt, and Warut Suksompong. 2016. The impossibility of extending random dictatorship to weak preferences. *Economics Letters* 141, C (2016), 44–47.

[8] Eric Budish. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.

[9] George Christodoulou, Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, Paul W Goldberg, Jie Zhang, and Jinshan Zhang. 2016. Social welfare in one-sided matching mechanisms. In *International Conference on Autonomous Agents and Multiagent Systems*. 30–50.

[10] Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, and Jie Zhang. 2014. Social welfare in one-sided matchings: Random priority and beyond. In *International Symposium on Algorithmic Game Theory*. Springer, 1–12.

[11] Michael R. Garey and David S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman.

[12] Guangda Huzhang, Xin Huang, Shengyu Zhang, and Xiaohui Bei. 2017. Online Roommate Allocation Problem. In *International Joint Conference on Artificial Intelligence*. 235–241.

[13] Aanund Hylland and Richard Zeckhauser. 1979. The efficient allocation of individuals to positions. *Journal of Political Economy* 87, 2 (1979), 293–314.

[14] Piotr Krysta and Jinshan Zhang. 2016. House markets with matroid and knapsack constraints. In *International Colloquium on Automata, Languages, and Programming*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

[15] Van Bang Le and Florian Pfender. 2014. Complexity Results for Rainbow Matchings. *Theoretical Computer Science* 524, C (2014).

[16] Sagar Massand and Sunil Simon. 2019. Graphical one-sided markets. In *International Joint Conference on Artificial Intelligence*. 492–498.

[17] Lars-Gunnar Svensson. 1999. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare* 16, 4 (1999), 557–567.

A OMITTED PROOFS FROM SECTION 5

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 $\mathbb{E}(SW(A)) \geq \frac{1}{4}OPT(I)$.*

Proof. Consider the FF-CT-RSD * at iteration t; we define N^t be the left (unassigned) agents at time t, $\mathcal{G}^t$ be the topology unallocated plots at time t with vertex set $\mathcal{V}^t$ and edges $\mathcal{E}^t$, F^t be the singular friendship structure among N^t, OPT^t be the optimal allocation of N^t to the plot topology $\mathcal{G}^t$ and SW^t be the total social welfare by the algorithm by time t. We denote the past history of plot assignment before the t-th iteration as $\mathcal{H}^t$. We note that the $\text{SW}(\mathcal{A}) = \text{SW}^T$ (T be the total number of iterations) and $\text{SW}(\text{OPT}) = \text{SW}(\text{OPT}^0)$. Now, at each iteration t, we divide all friendship pairs into four sets as follows:

$$Z_1^t = \{(i,j) \in F^t : u_l(\text{OPT}^t) = 1 + \varphi; l = i, j\}$$
$$Z_2^t = \{(i,j) \in F^t : u_i(\text{OPT}^t) + u_j(\text{OPT}^t) = 2\varphi + 1\}$$
$$Z_3^t = \{(i,j) \in F^t : u_l(\text{OPT}^t) = \varphi; l = i, j\}$$
$$Z_4^t = \{(i,j) \in F^t : u_l(\text{OPT}^t) \leq 1; l = i, j\}$$

For convenience, we write $|Z_i^t| = z_i^t$ for all $i = 1, \ldots, 4$ and $|F^t| = f^t$. As friendship value $\varphi > 1$, at iteration t, if the first while loop runs then the randomly chosen pair will always be able to capture two adjacent plots due to higher friendship value φ than plot values and availability of adjacent plots. which implies that

$$\mathbb{E}[\text{SW}^{t+1}|\mathcal{H}^t] \geq \text{SW}^t + \frac{z_1^t}{f^t}(2\varphi+2) + \frac{z_2^t}{f^t}(2\varphi+1) + \frac{z_3^t + z_4^t}{f^t}(2\varphi) \quad (2)$$

Now, we analyze the decrease in the optimal welfare after each iteration when the first while loop condition satisfies ($\mathcal{E}^t \neq \emptyset$ and $F^t \neq \emptyset$). Consider randomly selected pair (i, j) at iteration t, they will grab any available adjacent plots among available plots $(v, w) \in \mathcal{E}^t$ due to higher friendship value than plot values. If randomly chosen $(i,j) \in Z_1^t \cup Z_2^t \cup Z_3^t$, then (i, j) can destroy at most 2 allocated adjacent plots to friends in OPT^t allocation and their own friendship. Therefor, it can cost at most 6φ in friendship value of the OPT^t allocation and if $(i, j) \in Z_4^t$ then it can cost at most 4φ in friendship value of the OPT^t allocation. Moreover, in the case when $(i,j) \in Z_1^t \cup Z_2^t \cup Z_4^t$, by grabbing adjacent plots (v, w), it can disturb at most 4 assigned plot (with value 1) in OPT^t allocation and when $(i, j) \in Z_3^t$, by grabbing adjacent plots (v, w), it can disturb at most 2 plot values in OPT^t allocation. Therefore we can write using equation 2;

$$\begin{aligned}&\mathbb{E}[\mathrm{SW}(\mathrm{OPT}^{t+1})-\mathrm{SW}(\mathrm{OPT}^{t})|\mathcal{H}^t]\\&\leq \frac{z_1^t+z_2^t}{f^t}(6\varphi+4)+\frac{z_3^t}{f^t}(6\varphi+2)+\frac{z_4^t}{f^t}(4\varphi+4)\\&\leq \frac{4(2\varphi+2)z_1^t}{f^t}+\frac{4(2\varphi+1)z_2^t}{f^t}+\frac{8\varphi(z_3^t+z_4^t)}{f^t}\\&\leq 4\cdot\mathbb{E}[\mathrm{SW}^{t+1}-\mathrm{SW}^t|\mathcal{H}^t]\end{aligned}$$

Once all pairs of friends or adjacent edges in the topology are exhausted ($F^t = \emptyset$ or $\mathcal{E}^t = \emptyset$), our algorithm becomes random serial dictatorship in one-sided matching market. Theorem 2 in [3] implies that

$$\begin{aligned}&\mathbb{E}[\mathrm{SW}(\mathrm{OPT}^{t+1})-\mathrm{SW}(\mathrm{OPT}^{t})|\mathcal{H}^t]\leq\\&3\cdot\mathbb{E}[\mathrm{SW}^{t+1}-\mathrm{SW}^t|\mathcal{H}^t]\leq 4\cdot\mathbb{E}[\mathrm{SW}^{t+1}-\mathrm{SW}^t|\mathcal{H}^t]\end{aligned}$$

This implies that the sequence of random variables $X^0 = 0$ and $X^t - X^{t-1} = 4\cdot(\mathrm{SW}^t-\mathrm{SW}^{t-1}) - (\mathrm{SW}(\mathrm{OPT})^t - \mathrm{SW}(\mathrm{OPT})^{t-1})$ is a sub-martingale as $\mathbb{E}[X^{t+1}|\mathcal{H}^t] \geq X^t$. Therefore by Doobs Stopping Theorem, we get $\mathbb{E}[X^T|\mathcal{H}^t] \geq 0$ which implies,

$$\begin{aligned}0&\leq\mathbb{E}\Big[\sum_{k=1}^{T}(X^k-X^{k-1})\Big]=4\cdot\mathbb{E}\Big[\sum_{k=1}^{T}(\mathrm{SW}^k-\mathrm{SW}^{k-1})\Big]\\&-\mathbb{E}\Big[\sum_{k=1}^{T}(\mathrm{SW}(\mathrm{OPT}^{k-1})-\mathrm{SW}(\mathrm{OPT}^k))\Big]\\&\Rightarrow\mathbb{E}[\mathrm{SW}(\mathcal{A})]\geq\frac{\mathrm{OPT}}{4}\end{aligned}$$

This concludes the proof. □

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

Proof. Let $N^t, \mathcal{G}^t, \mathcal{V}^t, \mathcal{E}^t, \mathcal{H}^t, \mathrm{SW}^t$ and OPT^t are defined similar to the proof of Theorem 5.2 at $t-$th iteration of On-CA-RSD *. We denote $|N^t|$ as n^t for simplification. We define the subsets of the N^t as follows:

$Z_1^t = \{i \in N^t : i \text{ has friend } j, u_l(\mathrm{OPT}^t) = \varphi+1; l = i, j\}$

$Z_2^t = \{i \in N^t : i \text{ has friend } j, u_i(\mathrm{OPT}^t) + u_j(\mathrm{OPT}^t) = 1+2\varphi\}$

$Z_3^t = \{i \in N^t : i \text{ has friend } j, u_l(\mathrm{OPT}^t) = \varphi; l = i, j\}$

$Z_4^t = \{i \in N^t : i \text{ has no friend}, u_l(\mathrm{OPT}^t) \leq 1; l = 1, 2\}$

$Z_5^t = \{i \in N^t : i \text{ has no friend}, u_i(v) = 0, \forall v \in \mathcal{V}^t\}$

Consider t with $\mathcal{E}^t \neq \emptyset$, If randomly chosen $i \in Z_1^t \cup Z_2^t \cup Z_3^t \cup Z_4^t$, as value of friendship $\varphi > 1$ and $\mathcal{E}^t \neq \emptyset$, i will declare her friend j and pick a plot $v \in \mathcal{V}^t$ with available adjacent plot with maximizing her own utility then if possible her friends utility. If $i \in N^t \setminus (Z_1^t \cup Z_2^t \cup Z_3^t \cup Z_4^t \cup Z_5^t)$, i will pick a plot with value 1. We note that the On-CA-RSD never picks $i \in Z_5^t$. This implies

$$\begin{aligned}\mathbb{E}[\mathrm{SW}^{t+1}|\mathcal{H}^t]\geq\mathrm{SW}^t+\frac{z_1^t}{n^t}(2\varphi+2)+\frac{z_2^t}{2}(2\varphi+1)\\+\frac{(z_3^t+z_4^t)2\varphi}{n^t}+1-\frac{(z_1^t+z_2^t+z_3^t+z_4^t+z_5^t)}{n^t}\end{aligned}\tag{3}$$

We now analyse the decrease in optimal welfare at iteration t whenever $\mathcal{E}^t \neq \varphi$. If the randomly chosen agent i at iteration t belongs to $Z_1^t \cup Z_2^t \cup Z_3^t$ then the agent i will pick the plot $v \in \mathcal{V}^t$ and force her friend j to pick $w \in \mathcal{N}_v$. Therefore in any case, by grabbing two adjacent plots $(v, w) \in \mathcal{E}^t$, they can destroy at most 2 other friendship values, their own friendship value in the OPT^t allocation where if $i \in Z_4^t$ then it can destroy at most 2 friendship values in OPT^t allocation (as they are not assigned adjacent plots in OPT^t allocation). Now, we analyse the decrease in plot values in optimal welfare. If $i \in Z_1^t \cup Z_4^t$, they can destroy at most 4 plot values in OPT^t allocation. Similarly, if $i \in Z_3^t \cup (N^t \setminus (Z_1^t \cup Z_2^t \cup Z_3^t \cup Z_4^t \cup Z_5^t))$, it can destroy at most 2 plot values in OPT^t allocation and for $i \in Z_2^t$ can destroy at most 3 plot values in OPT^t allocation (only one of them assigned to high valued plot in OPT^t). This implies that:

$$\begin{aligned}&\mathbb{E}[\mathrm{SW}(\mathrm{OPT}^{t})-\mathrm{SW}(\mathrm{OPT}^{t+1})|\mathcal{H}^t]\\&\leq\frac{z_1^t(6\varphi+4)}{n^t}+\frac{z_2^t(6\varphi+3)}{n^t}+\frac{z_3^t(6\varphi+2)}{n^t}+\frac{z_4^t(4\varphi+4)}{n^t}\\&\qquad+\Big(1-\frac{z_1^t+z_2^t+z_3^t+z_4^t+z_5^t}{n^t}\Big)(2\varphi+2)\\&\leq\frac{z_1^t(2\varphi+2)^2}{n^t}+\frac{z_2^t(2\varphi+1)(2\varphi+2)}{n^t}+\frac{z_3^t(2\varphi)(2\varphi+2)}{n^t}+\\&\frac{z_4^t(2\varphi)(2\varphi+2)}{n^t}+\Big(1-\frac{z_1^t+z_2^t+z_3^t+z_4^t+z_5^t}{n^t}\Big)(2\varphi+2)\\&\leq(2\varphi+2)\cdot\mathbb{E}[\mathrm{SW}^{t+1}-\mathrm{SW}^t|\mathcal{H}^t]\end{aligned}$$

Note that $z_1^t+z_2^t+z_3^t+z_4^t+z_5^t \leq n^t$. Once all adjacent edges in the topology are exhausted ($\mathcal{E}^t = \emptyset$), our algorithm becomes a random serial dictatorship in a one-sided matching market. Theorem 2 in [3] implies that

$$\begin{aligned}&\mathbb{E}[\mathrm{SW}(\mathrm{OPT}^{t+1})-\mathrm{SW}(\mathrm{OPT}^{t})|\mathcal{H}^t]\leq\\&3\cdot\mathbb{E}[\mathrm{SW}^{t+1}-\mathrm{SW}^t|\mathcal{H}^t]\leq(2\varphi+2)\cdot\mathbb{E}[\mathrm{SW}^{t+1}-\mathrm{SW}^t|\mathcal{H}^t]\end{aligned}$$

By similar argument as Theorem 5.2, we obtain

$$\mathbb{E}[\mathrm{SW}(\mathcal{A})]\geq\frac{\mathrm{OPT}}{2\varphi+2}$$

which concludes the proof. □

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

Proof. The proof of the theorem is similar to the Theorem 5.2. Consider the same terminologies which were defined in Theorem 5.2. First we analyse the gain in social welfare at $t-$th iteration when $F^t \neq \emptyset$ and $\mathcal{E}^t \neq \varphi$. We notice that the adjacent plots are available at $t-$th iteration. If randomly selected pair $(i, j) \in Z_1^t$ then both can obtain their maximum possible utility $1+\varphi$, however, it becomes little tricky when $(i, j) \notin Z_1^t$. If $(i, j) \in Z_2^t$, then increment in social welfare at $t-$th iteration is at least $\min\{2, 2\varphi+1\}-$either (i, j) can grab the assigned plots in OPT^t or they both grab their respective high valued plots which are non-adjacent. Similarly, when $(i, j) \in Z_3^t \cup Z_4^t$, (i, j) grabs two adjacent plots or two non-adjacent plots where at least one of them is getting high valued plots. Which implies;

$$\mathbb{E}[\mathrm{SW}^{t+1}|\mathcal{H}^t] \geq \mathrm{SW}^t + \frac{z_1^t(2\varphi+2)}{f^t} + \frac{z_2^t(\min\{2, 2\varphi+1\})}{f^t} + \frac{(z_3^t+z_4^t)(\min\{1, 2\varphi\})}{f^t}$$

The decrease in the optimal welfare after t−th iteration should be upper bounded by a similar quantity as Theorem 5.2. Therefore for $\varphi < 1$ we can write:

$$\begin{aligned}
&\mathbb{E}[\mathrm{SW}(\mathrm{OPT}^{t+1}) - \mathrm{SW}(\mathrm{OPT}^t)|\mathcal{H}^t]\\
&\leq \frac{z_1^t+z_2^t}{f^t}(6\varphi+4) + \frac{z_3^t}{f^t}(6\varphi+2) + \frac{z_4^t}{f^t}(4\varphi+4)\\
&\leq \frac{4\left(1+\frac{1}{\varphi}\right)(2\varphi+2)z_1^t}{f^t} + \frac{4\left(1+\frac{1}{\varphi}\right)(\min\{2,2\varphi+1\})z_2^t}{f^t}\\
&\quad + \frac{4\left(1+\frac{1}{\varphi}\right)(\min\{2,2\varphi+1\})(z_3^t+z_4^t)}{f^t}\\
&\leq 4\left(1+\frac{1}{\varphi}\right)\cdot \mathbb{E}[\mathrm{SW}^{t+1} - \mathrm{SW}^t|\mathcal{H}^t]
\end{aligned}$$

Now, by the similar analysis as Theorem 5.3, we obtain the desired result.

□

Theorem A.2. *Let A be the output of* On-CA-RSD* *on a binary instance I with $\varphi < 1$. Then $\mathbb{E}(SW(A)) \geq \frac{\varphi}{4\varphi+4} OPT(I)$.*

Proof. The proof of the theorem is similar to the Theorem 5.2 and Theorem A.1. □

A.1 Mixed Integer Program

In this section, we showcase a mixed integer program (MIP) formulation for the HA+X problem.

$$max: \sum_{i\in N} U_i(A) \tag{4}$$

$$\sum_{i\in N} a_{i,v} \leq 1, \qquad \forall v \in \mathcal{V} \tag{5}$$

$$\sum_{v\in \mathcal{V}} a_{i,v} \leq 1, \qquad \forall i \in N \tag{6}$$

$$u_i = \sum_{v\in V}[u_i(v)a_{i,v} + \varphi \sum_{f\in F_i}\sum_{v'\in \mathcal{N}_v} a_{f,v'}], \qquad \forall i \in N \tag{7}$$

$$a_{i,v} \in \{0,1\}, \forall i \in N, v \in \mathcal{V} \tag{8}$$

Equation 7 is the core of the MIP formulation, and it encodes the utility function by hard-coding the friends and nearby functions for each buyer and cottage. Equations 5 and 6 require that each buyer be matched to at most one cottage, and each cottage be matched to at most one buyer. And Equation 8 encodes the binary variable constraints for $a_{b,c}$.