

Pure Nash Equilibria in Online Fair Division

Martin Aleksandrov

UNSW Sydney

Data61, CSIRO

TU Berlin

martin.aleksandrov@data61.csiro.au

Toby Walsh

UNSW Sydney

Data61, CSIRO

TU Berlin

toby.walsh@data61.csiro.au

Abstract

We consider a fair division setting in which items arrive one by one and are allocated to agents via two existing mechanisms: LIKE and BALANCED LIKE. The LIKE mechanism is strategy-proof whereas the BALANCED LIKE mechanism is not. Whilst LIKE is strategy-proof, we show that it is not group strategy-proof. Indeed, our first main result is that *no* online mechanism is group strategy-proof. We then focus on pure Nash equilibria of these two mechanisms. Our second main result is that computing a pure Nash equilibrium is tractable for LIKE and intractable for BALANCED LIKE. Our third main result is that there could be multiple such profiles and counting them is also intractable even when we restrict our attention to equilibria with a specific property (e.g. envy-freeness).

1 Introduction

Fair division is a fundamental problem our society faces today. It is even more challenging when the fair division decisions are made online often with only partial knowledge. For example, a food bank allocates food as it is donated. As a second example, an organ matching program allocates newly donated organs of deceased patients to transplant recipients within a few hours of their arrival. We cannot wait for more items to arrive before allocating existing items to agents. We focus here on a simple online fair division model introduced in [Walsh, 2014] where a new item arrives at each time step.

Fair division of indivisible items without monetary transfers is a particularly difficult case because items cannot be divided to satisfy multiple agents and because we cannot use money to compensate agents receiving less utility. Competing agents are then likely to behave strategically, especially if they have some information about what items remain to be allocated. We study here such strategic behavior for two promising mechanisms for online fair division: LIKE and BALANCED LIKE [Walsh, 2015]. Since our fair division problem is online, we typically have some uncertainty about items yet to arrive. This limits strategic behavior. However, even with limited information, agents may still be able to act strategically. For instance, with LIKE strategic behavior depends just on the items to arrive and not their arrival order.

We look at these two simple mechanisms for multiple reasons. They satisfy many nice axiomatic properties [Aleksandrov *et al.*, 2015]. For example, they are randomized and this helps them achieve properties such as fairness, anonymity and equal treatment of equals. The LIKE mechanism is strategy-proof, envy-free ex ante but is not fair ex post. In particular, one agent can have unbounded envy for another or even not receive a single item. LIKE characterizes envy-freeness ex ante in general and is one of the few strategy-proof mechanisms. On the other hand, BALANCED LIKE bounds the envy and guarantees that agents will be allocated similar numbers of items but it is not strategy-proof. Neither mechanism is Pareto efficient in general but both become so with binary utilities. Moreover, both LIKE and BALANCED LIKE are the most competitive mechanisms possible from an egalitarian and worst-case perspective. We, therefore, turn attention to questions around their *Nash equilibria*.

Nash equilibrium is a game-theoretic concept commonly used to understand the behavior of competing agents. If there is a unique pure Nash equilibrium, the main focus is usually around computational questions. One may also be interested in what the “best” equilibrium of the game is, or whether a given pure strategy is played in any equilibrium [Gilboa and Zemel, 1989]. If there are multiple pure Nash equilibria, the focus is then shifted to counting questions. One may argue that it is impossible to have a good overview of all the Nash equilibria of a game if one cannot even count them. This, however, does not preclude us from making predictions about equilibria. In our work, we look at both types of questions. For example, we compute pure Nash equilibria of both LIKE and BALANCED LIKE as well as count them.

There are several reasons why we study Nash equilibria of the one-shot game rather than the subgame perfect Nash equilibria of the repeated game. From a practical perspective, these games are typically played as one-shot. It is impractical to query agents repeatedly as every item arrives so we often collect their preferences in advance. For example, hundreds of items arrive in the food bank problem and, for this reason, charities request food only a few times per day when they reveal their utilities.

In addition, we may force agents to play the one-shot game to reduce their strategic options. Nevertheless, the problem remains online because the items arrive over time, they need to be allocated promptly and as we do not know in advance

what items actually will be donated. Our results also suppose we know utilities. Again, this is often reasonable even though mechanisms like LIKE and BALANCED LIKE are ordinal. For instance, in the food bank problem, the utility of liked items might simply be their retail price. This is public information. In the organ matching problem, doctors have a function to compute the utility of a matching based on age and medical history. This is again public information.

Computing equilibria is more relevant for agent control when agents behave strategically. Counting equilibria is more relevant for chair control when the chair wants to enforce a Pareto efficient and envy-free equilibrium once the items start to arrive. To achieve this, we show that the chair might have to choose amongst exponentially many such equilibria. Counting complexity is also important if the chair wants to limit say the envy-freeness of the equilibrium outcome. For example, at Harvard Business School, students submit their sincere preferences for courses and a “proxy draft” mechanism allocates courses to them by using a Pareto efficient equilibrium strategy [Budish and Cantillon, 2012].

Our contributions: LIKE is strategy-proof whereas BALANCED LIKE is not. Our first result is that *no* mechanism is group strategy-proof. These results motivate why we look at competitive pure Nash equilibria for BALANCED LIKE and group pure Nash equilibria for LIKE. Our second result is that computing a pure Nash equilibrium is tractable for LIKE and intractable for BALANCED LIKE. Our third result is that there could be multiple such profiles and counting them is also intractable in the worst-case. Our intractability results have many interesting consequences. For example, we show that deciding whether an equilibrium with egalitarian welfare (i.e. minimum expected utility) of at least k is intractable. As a second example, we show that counting envy-free pure Nash equilibria is also computationally hard.

2 Preliminaries

An instance \mathcal{I} of an online allocation problem has (1) *agents* a_1, \dots, a_n , (2) *indivisible items* o_1, \dots, o_m , (3) *private utility* u_{ij} for each agent a_i and each item o_j and (4) *ordering* o of the items. We consider *binary* utilities and *general* rational utilities. We say that agent a_i *likes* item o_j if $u_{ij} > 0$.

Online setting: The allocation proceeds in rounds. At round j , we suppose that item o_j when each agent a_i places a rational *bid* v_{ij} for item o_j . A *mechanism* now allocates item o_j to an agent supposing *no* information about future items.

We consider the LIKE and BALANCED LIKE from [Aleksandrov *et al.*, 2015]. Each of them allocates item o_j uniformly at random to a *feasible* agent. With LIKE, agent a_i is feasible for item o_j if $v_{ij} > 0$. With BALANCED LIKE, agent a_i is feasible for item o_j if $v_{ij} > 0$ and a_i has fewest items amongst those agents that bid positively for item o_j .

Pure Nash equilibria: Let $v_i = (v_{i1}, \dots, v_{im})$ denote the vector of bids that agent a_i report for all items and $v_{i,-j}$ the sub-vector of v_i without bid v_{ij} . Also, let V denote the bid matrix (i.e. *bidding profile*) and V_{-i} the bid sub-matrix of V without v_i . Further, let $\bar{u}_i(v_i, V_{-i}, o)$ denote the *expected utility* of agent a_i for all items. We say that v_i^1 is a *strict best response* vector of agent a_i to V_{-i} if

$\bar{u}_i(v_i^1, V_{-i}, o) > \bar{u}_i(v_i^2, V_{-i}, o)$ for each vector $v_i^2 \neq v_i^1$. We say that v_i^1 is a *weak best response* vector of agent a_i to V_{-i} if $\bar{u}_i(v_i^1, V_{-i}, o) = \bar{u}_i(v_i^2, V_{-i}, o)$ for some vector $v_i^2 \neq v_i^1$ and $\bar{u}_i(v_i^1, V_{-i}, o) \geq \bar{u}_i(v_i^3, V_{-i}, o)$ for each vector $v_i^3 \neq v_i^1$. We next define pure Nash equilibria. A bidding matrix V is a *strict pure Nash equilibrium* if all bidding vectors in V are strict pure best responses, and V is a *weak pure Nash equilibrium* if some vectors in V are weak best responses and the other vectors in V are strict best responses. Throughout the paper, we consider *monotone* pure Nash equilibria in which each agent bids their positively or zero for each item they like, and bids zero for each item they dislike.

Agent a *envies (ex ante) ex post* agent b if a 's (expected) utility of b 's (expected) allocation is greater than a 's (expected) utility of a 's (expected) allocation. A bidding profile is *envy-free (ex ante) ex post* for a mechanism if no agent envies (ex ante) ex post another one given the (expected) allocation returned by the mechanism on this profile. A bidding profile is *(ex ante) ex post Pareto efficient* for a mechanism if the (expected) allocation returned by the mechanism on this profile is Pareto optimal. A mechanism is *strategy-proof* if, with complete information, no agent can misreport their sincere utilities for items and thus strictly increase their (expected) utility. A mechanism is *group strategy-proof* if, with complete information, no set (or no *group*) of agents can misreport their sincere utilities and thus strictly increase their group expected outcome (i.e. the sum of the expected utilities of the agents in the group).

We show our hardness results using notions from computational complexity and graph theory.

Computational complexity: We use complexity classes such as P, NP, coNP and #P as well as mappings such as *Karp*, *Turing* and *parsimonious* reductions [Garey and Johnson, 1979; Turing, 1936; Valiant, 1979].

Graph theory: Let G be a bipartite graph. A *matching* in G is a set of vertex-disjoint edges. We say that it *matches* a vertex if there is an edge in it that is incident with the vertex. A matching is *perfect* if it matches all vertices in G .

We report our results for strategy-proofness in Section 3. We then study how to compute equilibria in Section 4. In Section 5 and Section 6, we show non-uniqueness of equilibria and hardness of counting them. We conclude in Section 7.

3 Manipulable Mechanisms

We know that the LIKE mechanism is strategy-proof and the BALANCED LIKE mechanism is not strategy-proof [Aleksandrov *et al.*, 2015]. We, therefore, proceed with group manipulations. In our food allocation problem, charities in the same region might decide to bid together for all items they like in the allocation. In our organ matching problem, hospitals in the same city might decide to report common patients multiple times on a national level and thus bias the national matching of organs in their favor. Group strategy-proof mechanisms exist and are successfully characterized in many domains [Martínez *et al.*, 2004; Pountourakis and Vidali, 2012]. Unfortunately, in online fair division, no mechanism is group strategy-proof in general.

Theorem 1 *With general utilities, no mechanism is group strategy-proof.*

Proof. Consider two agents and one item. Suppose the item is not always allocated deterministically to the agent bidding most for it. Then, from a group strategy perspective, the two agents can improve the group outcome by having only the agent with greatest utility bidding for the item. Hence, to be group strategy-proof, the item must be deterministically allocated to agent bidding most for it. However, the underbidder can now improve their individual outcome by increasing their bid to exceed the current winner's bid. Thus, deterministically allocating item to agent bidding most for it is not strategy-proof and thus not group strategy-proof. \square

If we use LIKE in the proof of Theorem 1, we immediately conclude that this mechanism is not group strategy-proof even with 0/1/2 utilities. Surprisingly, however, it becomes group strategy-proof with 0/1 utilities.

Theorem 2 *With 0/1 utilities, the LIKE mechanism is group strategy-proof.*

Proof. Consider a group of agents. Note that the expected outcome of the group for each item o_j is independent of their bidding strategies for items o_1 to o_{j-1} . Each joint bidding strategy of the agents from the group can be represented as a sequence of individual agent strategies. There are multiple such sequences that correspond to a given joint bidding strategy. The group outcome increases with LIKE if the next two conditions hold. For each such sequence and each individual bidding strategy in it, we have that the group expected outcome does not decrease. For each such sequence and some individual bidding strategy in it, we have that the group expected outcome increases. We show that no such strictly increasing individual strategy could exist. Let us pick an agent from the group and an item, say a_1 and o_1 . WLOG, suppose that o_1 is liked by k agents from the group, and by l agents in total. If a_1 sincerely dislikes item o_1 and strategically bids 1 for it, l is in $[k, n-1]$ and the decrease from the sincere group outcome is $k/[l \cdot (l+1)]$. If a_1 sincerely likes item o_1 and strategically bids 0 for it, l is in $[k, n]$ and the decrease from the sincere group outcome is $(l-k)/[l \cdot (l-1)]$. Both deviations are non-negative. Hence, sincere bidding dominates. \square

By Theorem 2, it quickly follows that the sincere strategy is dominant for a group of agents and for an item if the agents from the group that like the item have the same utility for it.

Corollary 1 *With general utilities, the LIKE mechanism is group strategy-proof if, for each group G and item o_j , each agent from G that likes o_j has utility u_{G_j} for it.*

We next report one interesting observation about group strategy-proofness. Envy-freeness ex ante is incompatible with ex ante Pareto efficiency in general.

Observation 1 *With general utilities, no mechanism is group strategy-proof and envy-free ex ante, or group strategy-proof and ex ante Pareto efficient even with randomisation.*

The LIKE mechanism is group strategy-proof and envy-free ex ante with 0/1 utilities but not even bounded envy-free ex post. In contrast, BALANCED LIKE is envy-free ex ante and bounded envy-free ex post but not even strategy-proof. Both of them are Pareto efficient with 0/1 utilities.

4 Computing Pure Nash Equilibria

We show that there could be a unique weak or strict pure Nash equilibrium and then that computing one is intractable.

Example 1 (Unique weak pure Nash equilibrium) *Consider items o_1 and o_2 and agents a_1 and a_2 . Suppose that a_1 likes o_1 with 1 and o_2 with 2, and a_2 likes o_1 with 2 and o_2 with 1. Let us run the BALANCED LIKE mechanism. The sincere profile is not a pure Nash equilibrium. By bidding sincerely, each agent receives expected utility of $3/2$. By bidding strategically 0 for item o_1 , each agent receives strictly higher expected utility of 2. This is the only pure Nash equilibrium which happens to be weak because the agents receive the same outcomes if a_2 reports 0 for o_2 in it. \square*

Example 2 (Unique strict pure Nash equilibrium) *Consider the instance in the proof of Theorem 2 from [Aleksandrov et al., 2015]. It has a unique pure Nash equilibrium which happens to be strict. \square*

We further study the following computational problem for an instance and a mechanism.

PURE NASH EQUILIBRIUM
 Input: instance \mathcal{A} and mechanism \mathcal{M}
 Output: a pure Nash equilibrium of \mathcal{A} with \mathcal{M}

If we can solve this problem in polynomial time, then we can solve its *decision* version in polynomial time. The decision version asks whether there is a pure Nash equilibrium of the given instance with the given mechanism. We show that the decision version is intractable for BALANCED LIKE and therefore the computational problem is intractable as well. Our Reduction 1 is a Karp mapping from the *minimum size maximal matching problem* in 3-regular bipartite graphs which is NP-complete [Demange and Ekim, 2008].

Reduction 1 Let R be a 3-regular bipartite graph with vertices u_1, \dots, u_N and v_1, \dots, v_N . For each vertex u_i , let v_{i1}, v_{i2}, v_{i3} denote the vertices connected to it and $e_{3 \cdot (i-1)+1} = (u_i, v_{i1})$, $e_{3 \cdot (i-1)+2} = (u_i, v_{i2})$, $e_{3 \cdot (i-1)+3} = (u_i, v_{i3})$ the edges incident with it. Each edge e_k can be represented as (u_i, v_j) for some $u_i \in \{u_1, \dots, u_N\}$ and $v_j \in \{v_{i1}, v_{i2}, v_{i3}\}$. We construct instance \mathcal{I}_R as follows:

- **Agents:** 1 agent a_k per edge e_k , agents b and c and agents d_1 to $d_{2 \cdot (N-r)}$ (i.e. $5 \cdot N - 2 \cdot r + 2$ agents),
- **Items:** 1 item per vertex v_j , 2 items u_{i1} and u_{i2} per vertex u_i , 2 more items w_{i1} and w_{i2} per vertex u_i and items x_1 to x_{N-r} , y (i.e. $6 \cdot N - r + 1$ items),
- **Non-zero Utilities** each agent a_k has utility of 1 for items $v_j, u_{i1}, u_{i2}, w_{i1}, w_{i2}$ and x_1 to x_{N-r} ; agent b has utility of 1 for items x_{N-r} and y ; agent c has utility of 1 for item y only; each agent d_k has utility of 1 for u_{11}, u_{12} to u_{N1}, u_{N2} , and
- **Ordering:** $o = (u_{11}u_{12} \dots u_{N1}u_{N2}v_1 \dots v_Nw_{11}w_{12} \dots w_{N1}w_{N2}x_1 \dots x_{N-r}y)$.

The completely detailed proofs of the next technical Lemmas 1 and 2 are omitted for reasons of space.

Lemma 1 *There is a maximal matching in R of cardinality at most r iff, with the BALANCED LIKE mechanism, agent b receives item y in \mathcal{I}_R with probability greater than zero.*

Lemma 2 *With the BALANCED LIKE mechanism, the sincere profile of \mathcal{I}_R is a pure Nash equilibrium.*

Proof sketch. Suppose the sincere profile is not a pure Nash equilibrium. Consequently, there is an agent that receives strictly higher expected utility by reporting zeros for items they sincerely like supposing the other agent strategies are the sincere ones. We consider four cases. In the first one, suppose this is agent a_1 . There is a *symmetry* in the preferences of agents a_1, a_2, a_3 , and each other triplet of a 's agents. There is another *symmetry* in the preferences of agents a_i, a_j, a_k such that all like item v_h for each h . Based on these symmetries, we obtained that agent a_k is sincere iff agent a_i is sincere, and that agent a_1 is strictly sincere. In the second case, suppose this is agent b . They receive outcome of $p + (1 - p) \cdot (1/2)$ if they bid 1 for x_{N-r}, y , $1/2$ if they bid 1 only for y , $p > 0$ if they bid 1 only for x_{N-r} , and 0 if they do not bid for items. Sincerity dominates. In the third case, suppose this is agent c . They receive $p + (1 - p) \cdot (1/2)$ if they bid 1 for y , and 0 otherwise. Sincerity dominates. In the fourth case, suppose this is agent d_1 . There is a *symmetry* in the preferences of agents d_1 to $d_{2 \cdot (N-r)}$. Based on it, we obtained that agent d_k is sincere iff agent d_i is sincere, and that agent d_1 is strictly sincere. \square

Lemma 3 *With the BALANCED LIKE mechanism, there is at least one more pure Nash equilibria of \mathcal{I}_R besides the sincere profile iff agent b receives item y in \mathcal{I}_R with probability zero.*

Proof. If agent b receives item y in \mathcal{I}_R with probability 0, then let us substitute their sincere bid of 1 for item y with their strategic bid of 0 for this item. The new profile is another pure Nash equilibrium in which all agents receive exactly the same outcomes as in the sincere profile. If there is at least one more pure Nash equilibrium of \mathcal{I}_R besides the sincere profile, then it must be the case that agent b receives item y with zero probability in the sincere profile. This follows because each other agent is strictly sincere by Lemma 2. \square

There are many hardness implications under Turing reductions from Lemmas 1, 2 and 3. There is a strict pure Nash equilibrium of \mathcal{I}_R iff the sincere profile is a strict pure Nash equilibrium. There is a weak pure Nash equilibrium of \mathcal{I}_R iff the sincere profile is not a strict pure Nash equilibrium. Hence, there is a weak pure Nash equilibrium of \mathcal{I}_R iff there is not a strict pure Nash equilibrium of \mathcal{I}_R . Further, there is a pure Nash equilibrium of \mathcal{I}_R iff there is a weak or a strict pure Nash equilibrium of \mathcal{I}_R iff agent b receives item y with zero or positive probability. Further, there is a unique pure Nash equilibrium of \mathcal{I}_R iff the sincere profile is a strict pure Nash equilibrium. There is a pure Nash equilibrium of \mathcal{I}_R with egalitarian welfare at least 1 (achieved by agent c for item y) iff the sincere profile is a weak pure Nash equilibrium. Even more, there is a (weak) strict pure Nash equilibrium of \mathcal{I}_R in which agent b bids 1 for item y iff the sincere profile is a (weak) strict pure Nash equilibrium. These problems have indirect implications to the problem of finding a Nash equilibrium in general [Dickhaut and Kaplan, 1993; Porter *et al.*, 2008; Sandholm *et al.*, 2005].

Corollary 2 *With 0/1 utilities and the BALANCED LIKE mechanism, problem PURE NASH EQUILIBRIUM is NP-hard.*

We next study group pure Nash equilibria for the LIKE mechanism. We can view the allocation process with this mechanism as a repeated game [Tomala, 1998]. This follows from the fact that this mechanism is Markovian. In each round j of this repeated game, there is a group pure Nash equilibrium for the current item o_j . The group pure Nash equilibrium of a game played m rounds is a collection that has a group pure Nash equilibrium for each item.

Theorem 3 *With general utilities and the LIKE mechanism, problem PURE NASH EQUILIBRIUM is P.*

Proof. WLOG, we show that computing a group pure Nash equilibrium for say item o_j is tractable. At step 1, let us pick a group of agents G . WLOG, suppose that all agents from G (no agent that dislikes o_j can increase the group outcome of G by bidding positively for o_j), and k other agents not from G like item o_j . We need to compute a set $H \subseteq G$ such that $H = \arg \max_{J \subseteq G} [(\sum_{a_i \in J} u_{ij}) / (|J| + k)]$. Note that $\max_{J \subseteq G} [(\sum_{a_i \in J} u_{ij}) / (|J| + k)]$ is equal to $\max_{l \in [0, |G|]} \max_{J_l \subseteq G} [(\sum_{a_i \in J_l} u_{ij}) / (l + k)]$ where l agents are in J_l . But, then $\max_{J_l \subseteq G} [(\sum_{a_i \in J_l} u_{ij}) / (l + k)]$ is equal to $(\max_{J_l \subseteq G} \sum_{a_i \in J_l} u_{ij}) / (l + k)$. Finally, the value of $\max_{J_l \subseteq G} \sum_{a_i \in J_l} u_{ij}$ can be computed in $\mathcal{O}(n)$ space and time for each l . To do so, we just pick l different agents from G with the first l greatest utilities for item o_j . Hence, the set of agents H can be computed in $\mathcal{O}(n^2)$ time. This result holds for any value of k . The “best” strategy for the agents from G is that the agents from H bid positively for item o_j and the agents from $G \setminus H$ bid zero for o_j . At step 2, consider next the set of k agents not from G that like item o_j and another group F that contains at least one of these agents. Similarly, compute the “best” strategy of the agents from F that like item o_j and are not from G . Repeat step 2 with the remaining agents not from G and F but that like o_j till not more such agents are left. The “best” strategies converge to a group pure Nash equilibrium. Interestingly, even though there might be an exponential number of not necessarily disjoint groups of agents, our procedure terminates in $\mathcal{O}(n^3)$ space and time. Hence, a group pure Nash equilibrium can be computed in $\mathcal{O}(n^3 \cdot m)$ space and time. \square

5 Multiple Pure Nash Equilibria

In general, there could be multiple pure Nash equilibria. Some of these could be strict and some others could be weak. There also could be payoff equivalent or payoff different pure Nash equilibria. In fact, the number of weak pure Nash equilibria with equal outcomes could be exponential in the number of agents.

Example 3 (Weak pure Nash equilibria) *We consider the fair division of items o_1 to o_n between agents a_1 to a_n . Suppose that a_1 likes each item with utility 1 and each other agent a_i likes only o_i with utility 1. With BALANCED LIKE, the sincere bidding strategy is dominant for each agent. However, agent a_1 receives the same expected outcome of 1 if they bid 1 for o_1 with probability 1 and, for each of the other $n - 1$ items, they bid 1 with probability 1 or bid 0 with probability 1. Hence, there are 2^{n-1} weak pure Nash equilibria that give the same output as the sincere profile. \square*

Moreover, there could also be exponentially many strict pure Nash equilibria with unequal outcomes.

Example 4 (Strict pure Nash equilibria) We use n gadgets, each of 4 agents and 4 items, to construct instance \mathcal{A} with $4n$ agents and $4n$ items that have 2^n different strict pure Nash equilibria. The i th gadget has agents a_{i1} to a_{i4} and items o_{i1} to o_{i4} . Agent a_{i1} likes items o_{i1} and o_{i4} with 1, item o_{i2} with $1 + 2i \cdot \epsilon$ and dislikes item o_{i3} . Agent a_{i2} likes only item o_{i2} with 1. Agent a_{i3} likes items o_{i1} and o_{i4} with 1, item o_{i3} with $1 + (2i + 1) \cdot \epsilon$ and dislikes item o_{i2} . Agent a_{i4} likes only item o_{i3} with 1. The items are revealed in ordering $o_i = (o_{i1}o_{i2}o_{i3}o_{i4})$. Suppose that $\epsilon < \frac{1}{4 \cdot n}$ and we use BALANCED LIKE mechanism. There are two strict pure Nash equilibria of this gadget if we consider it as an online instance by itself: (1) agent a_{i1} bids 0 for item o_{i1} and all other agents are sincere, and (2) agent a_{i3} bids 0 for item o_{i1} and all other agents are sincere. These two profiles are different as their expected allocations are different. The instance \mathcal{A} contains n such gadgets: for $i \in [1, n]$, agents a_{i1} to a_{i4} and items o_{i1} to o_{i4} . The entire ordering of \mathcal{A} is $(o_1 \dots o_n)$. The two strict pure Nash equilibria of each gadget are independent and different than the two strict pure Nash equilibria of each other gadget. Hence, \mathcal{A} admits 2^n different equilibria. \square

6 Counting Pure Nash Equilibria

Computing pure Nash equilibria is intractable in general. Even if we compute a single pure Nash equilibrium, we cannot be sure that the agents will play according to it simply because there might be multiple pure Nash equilibria. As a result, the task of counting (or enumerating) pure Nash equilibria becomes more relevant than simply computing such profiles. We hence show that counting equilibria (and therefore enumerating them) is intractable even when computing a single pure Nash equilibrium is tractable.

In particular, we ask questions around counting pure Nash equilibria with some specific property. For example, how many profiles are there whose outcome is equal to the outcome of the *sincere* profile or to the one of some *equilibrium*? Or, how many equilibria are Pareto efficient or envy-free? Further, strategic behavior might increase or decrease the welfare value [Aleksandrov *et al.*, 2015]. But, how many equilibria maximize the *egalitarian* welfare \bar{e} , the *utilitarian* welfare \bar{u} or the *Nash* \bar{n} welfare (i.e. the minimum expected utility of an agent, the sum or the product of the expected utilities of the agents, respectively)?

We give a single parsimonious reduction from the popular $\#P$ -complete *counting perfect matchings problem* that output the number of perfect matchings in a given undirected bipartite graph G . Counting perfect matchings problem is shown to be $\#P$ -hard on 3-regular bipartite graphs in [Dagum and Luby, 1992]. Our Reduction 2 is very simple but also insightful because it provides a very tight bound on the complexity of counting, parameterizing or approximating pure Nash equilibria (i.e. as many items as agents, 0/1 utilities, each agent likes 3 items, each item is liked by 3 agents, each pair of agents like at most 2 items in common, the ordering is fixed, etc.) [Jerrum and Sinclair, 1989; Downey and Fellows, 2013].

Reduction 2 Let R be a 3-regular bipartite graph with vertices u_1, \dots, u_N and v_1, \dots, v_N . For each vertex u_i , let v_{i1}, v_{i2}, v_{i3} denote the vertices connected to it and $e_{3 \cdot (i-1)+1} = (u_i, v_{i1})$, $e_{3 \cdot (i-1)+2} = (u_i, v_{i2})$, $e_{3 \cdot (i-1)+3} = (u_i, v_{i3})$ the edges incident with it. Each edge e_k can be represented as (u_i, v_j) for some $u_i \in \{u_1, \dots, u_N\}$ and $v_j \in \{v_{i1}, v_{i2}, v_{i3}\}$. The allocation instance \mathcal{E}_R is as follows:

- **Agents:** 1 agent a_k per edge e_k (i.e. $3 \cdot N$ agents),
- **Items:** 1 item per vertex v_j , 2 items u_{i1}, u_{i2} per vertex u_i (i.e. $3 \cdot N$ items),
- **Non-zero Utilities** each agent a_k has utility of 1 for items v_j, u_{i1} , and u_{i2} , and
- **Ordering:** $o = (v_1 \dots v_N u_{11} u_{12} \dots u_{N1} u_{N2})$.

Theorem 4 Even with 0/1 utilities and BALANCED LIKE, counting pure Nash equilibria is $\#P$ -hard.

Proof. In the sincere profile, it is easy to show that the expected utility of each agent is equal to 1. In fact, the sincere profile in \mathcal{E}_R is a pure Nash equilibrium. Let us next consider the following *perfect* profile for each perfect matching in R . WLOG, for each $i \in [1, N]$, suppose the matching matches u_i to v_j which happens to be v_{i1} . The perfect profile that corresponds to this matching is as follows: for $i \in [1, N]$, agent $a_{3 \cdot (i-1)+1}$ bids 1 for items v_j, u_{i1}, u_{i2} , and agents $a_{3 \cdot (i-1)+2}$ and $a_{3 \cdot (i-1)+3}$ bid 1 only for items u_{i1}, u_{i2} . There is 1-to-1 correspondence between the perfect matchings in R and the perfect profiles in \mathcal{E}_R . Therefore, counting such profile is intractable. Note that computing one of them in \mathcal{E}_R is tractable [Hopcroft and Karp, 1973]. Finally, it is easy to show that each perfect profile is a pure Nash equilibrium in which each agent receives expected utility of 1. \square

Corollary 3 With BALANCED LIKE, counting pure Nash equilibria that maximize the egalitarian, utilitarian or Nash welfare is $\#P$ -hard.

Proof. For \bar{e} , each agent in each perfect profile of \mathcal{E}_R receives expected utility of 1. Therefore, the value of \bar{e} in this profile is 1. We argue that this is the maximum value of this welfare. To see this, consider any bidding profile in which some agents receive expected utility more than 1. Hence, there are some other agents that receive expected utility less than 1. This is guaranteed because there are as many items as agents in \mathcal{E}_R , the utilities are binary, each agent likes at least one item and each item is liked by at least one agent. As a result, the sum of the expected utilities of the agents is equal to n and hence each perfect profile maximizes \bar{e} . For \bar{u} , each perfect profile maximizes it because at least one agent bids for each item. For \bar{n} , we argue that the maximum value of this welfare is 1 and is achieved whenever each agent receives expected outcome of 1. We highlight the key idea. Consider a triplet of agents, say $a_{3 \cdot (i-1)+1}$, $a_{3 \cdot (i-1)+2}$ and $a_{3 \cdot (i-1)+3}$, and a bidding profile in which say $a_{3 \cdot (i-1)+1}$ receives expected utility greater than 1. We showed that $a_{3 \cdot (i-1)+2}$ or $a_{3 \cdot (i-1)+3}$ receives expected utility smaller than 1 in such a profile. The same holds for every triplet of agents. For each such bidding profile, we obtained that $\bar{n} < 1$. Finally, each agent receives outcome of 1 in all profiles that $\bar{n} = 1$. \square

Corollary 4 With BALANCED LIKE, counting pure Nash equilibria whose outcome is sincere is $\#P$ -hard.

Proof. The result follows immediately because agents in \mathcal{E}_R receive expected outcomes of 1 in the sincere profile and in each perfect profile. \square

Corollary 5 *With BALANCED LIKE, counting pure bidding profiles whose outcome is equilibrium is #P-hard.*

Proof. Consider the following profile for each perfect profile. For $i \in [1, N]$, agent $a_{3 \cdot (i-1)+1}$ bids 1 for item v_j , agent $a_{3 \cdot (i-1)+2}$ bids 1 for item u_{i1} and agent $a_{3 \cdot (i-1)+3}$ bids 1 for item u_{i2} . We call this profile *deterministic* because a single agent bids for each item. Each agent receives (expected) utility of 1 in such a deterministic profile. This outcome is equal to their outcome in the sincere profile and in each perfect profile. However, the deterministic profile is not a pure Nash equilibria. Agent $a_{3 \cdot (i-1)+3}$ increases their outcome from 1 to $\frac{3}{2}$ by bidding 1 instead of 0 for item u_{i1} supposing we keep the other bidding strategies fixed. \square

Corollary 6 *With BALANCED LIKE, counting pure Nash equilibria whose outcome is Pareto efficient is #P-hard.*

Proof. Each perfect profile of \mathcal{E}_R is Pareto efficient ex post and ex ante. These profiles are pure Nash equilibria. \square

Corollary 7 *With BALANCED LIKE, counting pure Nash equilibria whose outcome is envy-free is #P-hard.*

Proof. Each perfect profile of \mathcal{E}_R is envy-free ex ante because each agent receives expected utility of 1. Each such profile is also envy-free ex post because each agent receives exactly one item in each possible allocation. These profiles are pure Nash equilibria. \square

Counting competitive pure Nash equilibria that are Pareto efficient or envy-free or maximize the welfares is intractable. We next count group pure Nash equilibria with these mechanisms. By Theorem 4, we conclude that counting group pure Nash equilibria is #P-hard with BALANCED LIKE even when each agent is alone in a group. This counting problem remains intractable for LIKE even with 0/1/2 utilities.

Reduction 3 *Given the 3-regular bipartite graph R in Reduction 2, construct an online allocation instance \mathcal{F}_R with the same set of agents, the same set of items, the same zero utilities and the same ordering as in \mathcal{E}_R . Only the non-zero utilities in \mathcal{F}_R differ from the ones in \mathcal{E}_R . For i between 1 and N and $j \in \{1, 2, 3\}$, agent $a_{3 \cdot (i-1)+j}$ has utility 2 for item v_{ij} if the vertices u_i and v_{ij} in the graph are matched in some perfect matching in R and 1 otherwise, and utility 2 for items u_{i1} and u_{i2} . To compute such a “coloring” in polynomial time, we can use a polynomial maximum matching algorithm such as the one from [Hopcroft and Karp, 1973].*

Theorem 5 *Even with 0/1/2 utilities and LIKE, counting group pure Nash equilibria is #P-hard.*

Proof. Consider instance \mathcal{F}_R and suppose that all agents are in a group. The sincere profile of \mathcal{F}_R is not a pure Nash equilibrium. Suppose sincere play. Each agent receives expected utility 2 if they have cardinal utility 2 for each of the three items they like, and $5/3$ if they have cardinal utility 1 for the first item they like and 2 for each of the remaining two items. Suppose next strategic play. Let the agents bid according to

one perfect profile that is defined for \mathcal{F}_R as for \mathcal{E}_R in Theorem 4. Each agent receives expected utility 2 given such a profile. Hence, the group outcome can only increase. In fact, such a perfect profile is a group pure Nash equilibrium. \square

Finally, note that each perfect profile of \mathcal{F}_R in the proof of Theorem 5 is also Pareto efficient and envy-free whereas the sincere profile is not. We conclude the next immediate result for LIKE (or even for BALANCED LIKE).

Corollary 8 *With LIKE, counting group pure Nash equilibria whose outcome is Pareto efficient, envy-free or maximizes the egalitarian, the utilitarian or the Nash welfare is #P-hard.*

7 Related Work and Conclusions

Our work is in-line with many other research works that study equilibria induced by allocation mechanisms. For example, the equilibria of the popular probabilistic serial rule are well-understood [Aziz *et al.*, 2015]. We consider weak and strict Nash equilibria as in [Shapley and Rigby, 1959] but we restricted ourselves to pure such profiles. The presentation of our counting results was inspired by the hardness results for two agents presented in [Conitzer and Sandholm, 2008]. Unfortunately, we could not inherit any of these results because manipulations for two agents are tractable for LIKE and BALANCED LIKE. Our decision problems are similar to other such problems in stochastic game-theory [Ummels and Wojtczak, 2009].

To conclude, we studied pure Nash equilibria for two allocation mechanisms in a simple online fair division model: LIKE and BALANCED LIKE. The LIKE mechanism is strategy-proof but not group strategy-proof and therefore we studied its group pure Nash equilibria. These profiles always exist and are tractable. But, are they always Pareto efficient or envy-free? The BALANCED LIKE mechanism is not strategy-proof and therefore we studied its competitive pure Nash equilibria. These profiles are intractable. It is an interesting future work to consider in more details the case of two agents. Even more, we want to explore possible parameterizations and approximations of these problems. For example, a sampling scheme might exist for our counting problems whose convergence is guaranteed by the Hoeffding’s inequality [Sabán and Sethuraman, 2013]. Finally, as our first next step, we plan to study the induced subgame perfect Nash equilibria of the repeated game. We also left *open* the question whether there is always a pure Nash equilibrium in such games for BALANCED LIKE.

Acknowledgements

Data61 (formerly NICTA) is supported by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program.

Toby Walsh acknowledges support from the European Research Council, as well as the Asian Office of Aerospace Research and Development.

References

- [Aleksandrov *et al.*, 2015] Martin Aleksandrov, Haris Aziz, Serge Gaspers, and Toby Walsh. Online fair division: Analysing a food bank problem. In *Proc. of IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 2540–2546, 2015.
- [Aziz *et al.*, 2015] Haris Aziz, Serge Gaspers, Simon Mackenzie, Nicholas Mattei, Nina Narodytska, and Toby Walsh. Equilibria under the probabilistic serial rule. In *Proc. of IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 1105–1112, 2015.
- [Budish and Cantillon, 2012] Eric Budish and Estelle Cantillon. The multi-unit assignment problem: Theory and evidence from course allocation at Harvard. *American Economic Review*, 102(5):2237–2271, 2012.
- [Conitzer and Sandholm, 2008] Vincent Conitzer and Tuomas Sandholm. New complexity results about nash equilibria. *Games and Economic Behavior*, 63(2):621–641, 2008.
- [Dagum and Luby, 1992] Paul Dagum and Michael Luby. Approximating the permanent of graphs with large factors. *Theoretical Computer Science*, 102(2):283–305, 1992.
- [Demange and Ekim, 2008] Marc Demange and Tinaz Ekim. Minimum maximal matching is NP-hard in regular bipartite graphs. In *5th International TAMC 2008, Xi'an, China, April 25-29, 2008. Proceedings*, pages 364–374, 2008.
- [Dickhaut and Kaplan, 1993] John Dickhaut and Todd Kaplan. *A Program for Finding Nash Equilibria*, pages 148–166. Springer New York, New York, NY, 1993.
- [Downey and Fellows, 2013] Rodney G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, 2013.
- [Garey and Johnson, 1979] M. R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- [Gilboa and Zemel, 1989] Itzhak Gilboa and Eitan Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80–93, 1989.
- [Hopcroft and Karp, 1973] John E. Hopcroft and Richard M. Karp. An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. *SIAM J. of Comp.*, 2(4):225–231, 1973.
- [Jerrum and Sinclair, 1989] Mark Jerrum and Alistair Sinclair. Approximating the permanent. *SIAM J. of Comp.*, 18(6):1149–1178, 1989.
- [Martínez *et al.*, 2004] Ruth Martínez, Jordi Massó, Alejandro Neme, and Jorge Oviedo. On group strategy-proof mechanisms for a many-to-one matching model. *Int. J. Game Theory*, 33(1):115–128, 2004.
- [Porter *et al.*, 2008] Ryan Porter, Eugene Nudelman, and Yoav Shoham. Simple search methods for finding a nash equilibrium. *Games and Economic Behavior*, 63(2):642–662, 2008.
- [Pountourakis and Vidali, 2012] Emmanouil Pountourakis and Angelina Vidali. A complete characterization of group-strategyproof mechanisms of cost-sharing. *Algorithmica*, 63(4):831–860, 2012.
- [Sabán and Sethuraman, 2013] Daniela Sabán and Jay Sethuraman. The complexity of computing the random priority allocation matrix. In *Proc. of WINE 2013, Cambridge, MA, USA, December 11-14*, page 421, 2013.
- [Sandholm *et al.*, 2005] Tuomas Sandholm, Andrew Gilpin, and Vincent Conitzer. Mixed-integer programming methods for finding nash equilibria. In *Proceedings, The Twentieth National Conference on Artificial Intelligence and the Seventeenth Innovative Applications of Artificial Intelligence Conference, July 9-13, 2005, Pittsburgh, Pennsylvania, USA*, pages 495–501, 2005.
- [Shapley and Rigby, 1959] L. S. Shapley and Fred D. Rigby. Equilibrium points in games with vector payoffs. *Naval Research Logistics Quarterly*, 6(1):57–61, 1959.
- [Tomala, 1998] Tristan Tomala. Pure equilibria of repeated games with public observation. *International Journal of Game Theory*, 27(1):93–109, 1998.
- [Turing, 1936] Alan M. Turing. On computable numbers, with an application to the Entscheidungsproblem. *J. Proc. of the London Math. Soc.*, 2(42):230–265, 1936.
- [Ummels and Wojtczak, 2009] Michael Ummels and Dominik Wojtczak. Decision problems for nash equilibria in stochastic games. In *Computer Science Logic, 23rd International Workshop, CSL 2009, 18th Annual Conference of the EACSL, Coimbra, Portugal, September 7-11, 2009. Proceedings*, pages 515–529, 2009.
- [Valiant, 1979] Leslie G. Valiant. The complexity of computing the permanent. *Theor. Comp. Sc.*, 8:189–201, 1979.
- [Walsh, 2014] Toby Walsh. Allocation in practice. In *Proc. of 37th KI 2014, Stuttgart, Germany, September 22-26, 2014*, pages 13–24, 2014.
- [Walsh, 2015] Toby Walsh. Challenges in resource and cost allocation. In *Proc. of 29th AAAI 2015, Austin, Texas, USA, January 25-30*, pages 4073–4077, 2015.