

The Adjusted Winner Procedure: Characterizations and Equilibria

Haris Aziz

NICTA and UNSW, Sydney, Australia
haris.aziz@nicta.com.au

Simina Brânzei

Aarhus University, Denmark
simina@cs.au.dk

Aris Filos-Ratsikas

Aarhus University, Denmark
filosra@cs.au.dk

Søren Kristoffer Stiil Frederiksen

Aarhus University, Denmark
ssf@cs.au.dk

Abstract

The Adjusted Winner procedure is an important mechanism proposed by Brams and Taylor for fairly allocating goods between two agents. It has been used in practice for divorce settlements and analyzing political disputes. Assuming truthful declaration of the valuations, it computes an allocation that is *envy-free*, *equitable* and *Pareto optimal*. We show that Adjusted Winner admits several elegant characterizations, which further shed light on the outcomes reached with strategic agents. We find that the procedure may not admit pure Nash equilibria in either the discrete or continuous variants, but is guaranteed to have ϵ -Nash equilibria for each $\epsilon > 0$. Moreover, under *informed* tie-breaking, exact pure Nash equilibria always exist, are Pareto optimal, and their social welfare is at least $3/4$ of the optimal.

1 Introduction

How should one fairly allocate resources among multiple economic agents? The question of fair division is as old as civil society itself [Moulin, 2003], with recorded instances of the problem dating back to thousands of years ago¹. Fair division has been studied in an extensive body of literature in economics, mathematics, political science [Moulin, 2003; Robertson and Webb, 1998; Brams and Taylor, 1996a; Young, 1994], and more recently, computer science, as the fair allocation of resources is arguably relevant to the design of multi-agent systems [Chevaleyre *et al.*, 2006; Procaccia, 2013]. For example, in shared computing environments, resources such as CPU and memory get multiplexed such that each user can use their computing unit at their own pace and without concern for the activity of others accessing the system².

¹For example, Hesiod informally describes in *The Theogony* (circa 700 B.C.) a fair division protocol known today as “Cut-and-Choose”.

²This problem was stated by Fernando Corbato (1962) in the context of developing time-sharing operating systems.

The *Adjusted Winner* procedure was introduced by Brams and Taylor (1996b) as a highly desirable mechanism for allocating *multiple divisible resources* among two parties. The procedure requires the participants to declare their preferences over the items and the outcome satisfies strong fairness and efficiency properties. Adjusted Winner has been advocated as a fair division rule for divorce settlements [Brams and Taylor, 1996b], international border conflicts [Taylor and Pacelli, 2008], political issues [Denoon and Brams, 1997; Massoud, 2000], real estate disputes [Levy, 1999], water disputes [Madani, 2010], deciding debate formats [Lax, 1999] and various negotiation settings [Brams and Taylor, 2000; Raith, 2000]. For example, it has been shown that the agreement reached during Jimmy Carter’s presidency between Israel and Egypt is very close to what Adjusted Winner would have predicted [Brams and Togman, 1996]. Adjusted Winner has been patented by New York University and licensed to the law firm Fair Outcomes, Inc [Karp *et al.*, 2014].

Although the merits of Adjusted Winner have been discussed in a large body of literature, the procedure is still not fully understood theoretically. We provide two novel characterizations, together with an alternative interpretation that turns out to be very useful for analyzing the procedure.

In addition, as observed already in [Brams and Taylor, 1996a], the procedure is susceptible to manipulation. However, fairness and efficiency are only guaranteed when the participants declare their preferences honestly. In a review of a well-known book on Adjusted Winner by Brams and Taylor [2000], Nalebuff [2001] highlights the need for research in this direction:

...thus we have to hypothesize how they (the players) would have played the game and where they would have ended up.

In this paper, we answer these questions by studying the existence, structure, and properties of pure Nash equilibria of the procedure. Until now, our understanding of the strategic aspects has been limited to the case of two items [Brams and Taylor, 1996a] and experimental predictions [Daniel and Parco, 2005]; our work identifies conditions under which Nash equilibria exist and provides theoretical guarantees for the performance of the procedure in equilibrium.

<i>Continuous Procedure</i>	<i>Lexicographic tie-breaking</i>	<i>Informed tie-breaking</i>
pure Nash	✗	✓
ϵ -Nash	✓	✓
<i>Discrete Procedure</i>	<i>Lexicographic tie-breaking</i>	<i>Informed tie-breaking</i>
pure Nash	✗	✓
ϵ -Nash	✓(*)	✓

Table 1: Existence of pure Nash equilibria in Adjusted Winner. The (*) result holds when the number of points is chosen appropriately.

1.1 Our contributions

We start by presenting the first characterizations of Adjusted Winner. We show that among all protocols that split at most one item, it is the only one that satisfies Pareto-efficiency and equitability. Under the same condition, we further show that it is equivalent to the protocol that always outputs a maxmin allocation.

Next, we obtain a complete picture for the existence of pure Nash equilibria in Adjusted Winner. We find the following: neither the discrete nor the continuous variants of the procedure are guaranteed to have pure Nash equilibria, but they do have ϵ -Nash equilibria, for every $\epsilon > 0$. Additionally, under *informed* tie-breaking, pure Nash equilibria always exist for both variants of the procedure.

Finally, we prove that the pure Nash equilibria of Adjusted Winner are envy-free and Pareto optimal with respect to the true valuations and that their social welfare is at least 3/4 of the optimal.

Our results concerning the existence or non-existence of pure Nash equilibria are summarized in Table 1.

2 Background

We begin by introducing the classical fair division model for which the Adjusted Winner procedure was developed [Brams and Taylor, 1996a]. Let there be two agents, Alice and Bob, that are trying to split a set $M = \{1, \dots, m\}$ of divisible items. The agents have preferences over the items given by numerical values that express their level of satisfaction. Formally, let $\mathbf{a} = (a_1, a_2, \dots, a_m)$ and $\mathbf{b} = (b_1, \dots, b_m)$ denote their *valuation vectors*, where a_j and b_j are the values assigned by Alice and Bob to item j , respectively.

An *allocation* $W = (W_A, W_B)$ is an assignment of fractions of items (or *bundles*) to the agents, where $W_A = (w_A^1, \dots, w_A^m) \in [0, 1]^m$ and $W_B = (w_B^1, \dots, w_B^m) \in [0, 1]^m$ are the allocations of Alice and Bob, respectively.

The agents have additive *utility* over the items. Alice’s utility for a bundle W_A , given that her valuation is \mathbf{a} , is: $u_{\mathbf{a}}(W_A) = \sum_{j \in M} a_j \cdot w_A^j$. Bob’s utility is defined similarly. The agents are weighted equally, such that their utility for receiving all the resources is the same: $\sum_{i \in M} a_i = \sum_{i \in M} b_i$.

There are two main settings studied in this context: *discrete* and *continuous* valuations. In the discrete setting, valuations are positive natural numbers that add up to some integer P

and can be interpreted as *points* (or coins of equal size) that the agents use to acquire the items. For ease of notation, we will consider the equivalent interpretation of valuations as rationals with common denominator P , where the valuations sum to 1. In the *continuous* setting, the valuations are positive real numbers, which are without loss of generality normalized to sum to 1. These normalizations make procedures invariant to any rescaling of the bids [Karp *et al.*, 2014; Brams *et al.*, 2012].

2.1 The Adjusted Winner Procedure

The Adjusted Winner procedure works as follows. Alice and Bob are asked by a mediator to state their valuations \mathbf{a} and \mathbf{b} , after which the next two phases are executed.

Phase 1: For every item i , if $a_i > b_i$ then give the item to Alice; otherwise give it to Bob. The resulting allocation is (W_A, W_B) and without loss of generality, $u_{\mathbf{a}}(W_A) \geq u_{\mathbf{b}}(W_B)$.

Phase 2: Order the items won by Alice increasingly by the ratio a_i/b_i : $\frac{a_{k_1}}{b_{k_1}} \leq \dots \leq \frac{a_{k_r}}{b_{k_r}}$. From left to right, continuously transfer fractions of items from Alice to Bob, until an allocation (W'_A, W'_B) where both agents have the same utility is produced: $u_{\mathbf{a}}(W'_A) = u_{\mathbf{b}}(W'_B)$.

Let $AW(\mathbf{a}, \mathbf{b})$ denote the allocation produced by Adjusted Winner on inputs (\mathbf{a}, \mathbf{b}) , where $AW_A(\mathbf{a}, \mathbf{b})$ and $AW_B(\mathbf{a}, \mathbf{b})$ are the bundles received by Alice and Bob. Note that the procedure is defined for strictly positive valuations, so the ratios are finite and strictly positive numbers. Examples can be found on the Adjusted Winner website³ as well as in [Brams and Taylor, 1996a].

Adjusted Winner produces allocations that are *envy-free*, *equitable*, *Pareto optimal*, and *minimally fractional*. An allocation W is said to be Pareto optimal if there is no other allocation that strictly improves one agent’s utility without degrading the other agent. Allocation W is equitable if the utilities of the agents are equal: $u_{\mathbf{a}}(W_A) = u_{\mathbf{b}}(W_B)$, envy-free if no agent would prefer the other agent’s bundle, and minimally fractional if at most one item is split.

Envy-freeness of the procedure implies *proportionality*, where an allocation is proportional if each agent receives a bundle worth at least half of its utility for all the items. A procedure is called envy-free if it always outputs an envy-free allocation (similarly for the other properties).

3 Characterizations

In this section, we provide two characterizations of Adjusted Winner⁴ for both the discrete and continuous variants. We begin with a different interpretation of the procedure that is useful for analyzing its properties.

An allocation is *ordered* if it can be produced by sorting the items in decreasing order of the valuation ratios a_i/b_i

³<http://www.nyu.edu/projects/adjustedwinner/>.

⁴The results here refer to the case when the agents report their true valuations to the mediator. We discuss the strategic aspects of the procedure in Section 4.

and placing a boundary line somewhere (possibly splitting an item), such that Alice gets the entire bundle to the left of the line and Bob gets the remainder:

$$\underbrace{\frac{a_{k_1}}{b_{k_2}} \geq \frac{a_{k_2}}{b_{k_2}} \geq \dots \geq \frac{a_{k_i}}{b_{k_i}}}_{\text{Alice's allocation}} \geq \underbrace{\frac{a_{k_{i+1}}}{b_{k_{i+1}}} \geq \dots \geq \frac{a_{k_m}}{b_{k_m}}}_{\text{Bob's allocation}}$$

The placement of the boundary line could lead either to an integral or a minimally fractional allocation. Note that the allocation that gives all the items to Alice is also ordered (but admittedly unfair).

It is clear to see that Adjusted Winner produces an ordered allocation (using some tie-breaking rule for items with equal ratios) with the property that the boundary line is appropriately placed to guarantee equitability. This is the way we will be interpreting the procedure for the remainder of the paper. We start by characterizing Pareto optimal allocations.

Lemma 1 *For any valuations (\mathbf{a}, \mathbf{b}) and any tie-breaking rule, an allocation W is not Pareto optimal if and only if there exist items i and j such that Alice gets a non-zero fraction (possibly whole) of j , Bob gets a non-zero fraction (possibly whole) of i , and $a_i b_j > a_j b_i$.*

Proof: (\Leftarrow) If such items i, j exist, then consider the exchange in which Bob gives $\lambda_i > 0$ of item i to Alice and Alice gives $\lambda_j > 0$ of item j to Bob, where

$$\frac{b_i}{b_j} \lambda_i < \lambda_j < \frac{a_i}{a_j} \lambda_i$$

Since $a_i/a_j > b_i/b_j$, such λ_i and λ_j do exist. Then Alice's net change in utility is:

$$a_i \lambda_i - a_j \lambda_j > a_i \lambda_i - a_j \frac{a_i}{a_j} \lambda_i = 0,$$

while Bob's net change is:

$$b_j \lambda_j - b_i \lambda_i > b_j \lambda_j - b_i \left(\lambda_j \frac{b_j}{b_i} \right) > b_j \lambda_j - b_j \lambda_j = 0.$$

Thus the allocation is not Pareto optimal.

(\Rightarrow) If the allocation W is not Pareto optimal, then Alice and Bob can exchange positive fractions of items to get a Pareto improvement.

Consider such an exchange and let S_A be the set of items for which positive fractions are given by Alice to Bob. Let S_B be defined similarly for Bob. Without loss of generality, S_A and S_B are disjoint; otherwise we could just consider the net transfer of any items that are in both S_A and S_B . Let $j \in S_A$ be the item with the lowest ratio a_j/b_j , and $i \in S_B$ with the highest ratio a_i/b_i .

If $a_i b_j > a_j b_i$ then we are done. Otherwise, assume by contradiction that for each item $k \in S_A$ and $l \in S_B$ it holds $a_k b_l \geq a_l b_k$. Then $a_k/b_k \geq a_l/b_l$; but then any Pareto improving exchange involving the transfer of items from S_A and S_B is only possible if at least one agent gets a larger fraction of items without the other agent getting a smaller fraction, which is impossible. \square

By Lemma 1, a Pareto optimal allocation can be obtained by sorting the items by the ratios of the valuations and drawing a boundary line somewhere. No matter where the boundary line is, the allocation is Pareto optimal (even if not equitable); thus an allocation is Pareto optimal and splits at most one item if and only if it is ordered. From this we obtain our first characterization.

Theorem 1 *Adjusted Winner is the only Pareto optimal, equitable, and minimally fractional procedure. Any ordered equitable allocation can be produced by Adjusted Winner under some tie-breaking rule.*

Note that both Pareto optimality and equitability are necessary for the characterization. By restricting to Pareto optimal allocations only, then even the allocation that gives all the items to one agent is Pareto optimal, while by restricting to equitable allocations only, even an allocation that throws away all the items is equitable. Similarly when the agents have identical utilities for some items, then there exist Pareto optimal and equitable allocations that split more than one item. For example, if the two agents have identical utilities over all items, then the allocation that gives half of each item to each agent is equitable and Pareto optimal. However, in the case that the valuation are such that $a_i/b_i \neq a_j/b_j$ for all items $i \neq j$, then Adjusted Winner is exactly characterized by Pareto optimality and equitability.

Theorem 2 *If the valuations satisfy $a_i/b_i \neq a_j/b_j$ for all items $i \neq j$, then the only Pareto optimal and equitable allocation is the result of Adjusted Winner.*

Proof: Recall first that an allocation is maxmin if it maximizes the minimum utility of the agents. Notice that AW achieves the same level of utility for the agents. Now Assume there exists an allocation (α, β) that is Pareto optimal and equitable, but not a result of AW. Then the allocation is not ordered and there exist at least two items i and j such that both agents get a fraction of them. This contradicts Lemma 1, and so (α, β) does not exist. \square

An allocation is *maxmin* if it maximizes the minimum utility over both agents. From Lemma 3.3 [Dall'Aglio and Mosca, 2007], an allocation is maxmin if and only if it is Pareto optimal and equitable. Together with Theorem 1, this leads to another characterization.

Theorem 3 *Adjusted Winner is equivalent to the procedure that always outputs a maxmin and minimally fractional allocation.*

4 Equilibrium Existence

In this section, we study Adjusted Winner when the agents are *strategic*, that is, their reported valuations are not necessarily the same as their actual valuations. Let $\mathbf{x} = (x_1, x_2, \dots, x_m)$ and $\mathbf{y} = (y_1, y_2, \dots, y_m)$ be the *strategies* (i.e. declared valuations) of Alice and Bob respectively. Call (\mathbf{x}, \mathbf{y}) a *strategy profile*. We will refer to \mathbf{a} and \mathbf{b} as the *true values* of Alice and Bob. Note that since strategies are reported valuations they are positive numbers that sum to 1.

Since the input to Adjusted Winner is now a strategy profile (\mathbf{x}, \mathbf{y}) instead of (\mathbf{a}, \mathbf{b}) , this means that the properties of the procedure are only guaranteed to hold with respect to the declared valuations, and not necessarily the true ones⁵.

A strategy profile (\mathbf{x}, \mathbf{y}) is an ϵ -Nash equilibrium if no agent can increase its utility by more than ϵ by deviating to a different (pure) strategy. For $\epsilon = 0$, we obtain a *pure Nash equilibrium*.

The main result of this section is that ϵ -Nash equilibria always exist. Furthermore, using an appropriate rule for settling ties between items with equal ratios x_i/y_i , the procedure also has exact pure Nash equilibria. We start our investigations from simple tie-breaking rules.

The main result of this section is that Adjusted Winner is only guaranteed to have ϵ -Nash equilibria when $\epsilon > 0$ using standard tie-breaking. For the discrete case, this is achieved by the center setting the number of points or equivalently the denominator large enough. Furthermore, we prove that when using an appropriate rule for settling ties between items with equal ratios x_i/y_i , the procedure does admit pure Nash equilibria. We start our investigations from the standard tie-breaking rules.

4.1 Lexicographic Tie-Breaking

The classical formulation of Adjusted Winner resolves ties in an arbitrary deterministic way, for example by ordering the items lexicographically, such that items with lower indices come first.

Continuous Strategies

First, we consider the case of continuous strategies. We start with the following theorem.

Theorem 4 *Adjusted Winner with continuous strategies is not guaranteed to have pure Nash equilibria.*

Proof: Take an instance with two items and valuations (\mathbf{a}, \mathbf{b}) , where $b_1 > a_1 > a_2 > b_2 > 0$. Assume by contradiction there is a pure Nash equilibrium at strategies (\mathbf{x}, \mathbf{y}) , where $\mathbf{x} = (x, 1 - x)$ and $\mathbf{y} = (y, 1 - y)$. We study a few cases and show the agents can always improve.

Case 1: ($x \neq y$). W.l.o.g. $x > y$ (the case $x < y$ is similar). Then $\exists \delta \in \mathbb{R}$ with $x - \delta > y \Rightarrow 1 - x + \delta < 1 - y$, and Alice can improve by playing $\mathbf{x}' = (x - \delta, 1 - x + \delta)$, as the boundary line moves to the left of its former position.

Case 2: ($x = y < 1/2$). Here both agents report higher values on the item they like less; Alice's allocation is $(1, \lambda)$ while Bob's is $(0, 1 - \lambda)$, for some $\lambda \in (0, 1)$. Then $\exists \delta \in \mathbb{R}$ with $x + \delta < 1/2$. By playing $\mathbf{y}' = (x + \delta, 1 - x - \delta)$, Bob gets $(1, 1 - \lambda')$, for some $\lambda' \in (0, 1)$. This is a strict improvement since $a_1 > a_2$.

Case 3: ($x = y > 1/2$). Both agents report higher values on the item they like more. Bob gets $(1 - \frac{1}{2x}, 1)$ and Alice gets $(\frac{1}{2x}, 0)$, with utilities $u_{\mathbf{a}}(AW(\mathbf{x}, \mathbf{y})) = \frac{a_1}{2x}$ and $u_{\mathbf{b}}(AW(\mathbf{x}, \mathbf{y})) = (1 - \frac{1}{2x})b_1 + b_2$. Let $\delta \in (0, \min(1 -$

⁵We will show that in the equilibrium, the procedure guarantees some of the properties with respect to the true values as well.

$x, 2x - 1)$ such that $\delta < \max\left\{\frac{4x(x-a_1)}{2x-a_1}, \frac{4x(b_1-x)}{2x-b_1}\right\}$. Observe that since $b_1 > a_1$ and $2x - a_1$ and $2x - b_1$ are positive, at least one of $x - a_1$ and $b_1 - x$ is strictly positive and by continuity of the strategy space, such a δ exists. Now consider alternative profiles $(\mathbf{x}', \mathbf{y}) = ((x - \delta, 1 - x + \delta), (x, 1 - x))$ and $(\mathbf{x}, \mathbf{y}') = ((x, 1 - x), (x + \delta, 1 - x - \delta))$. Since $\delta < 2x - 1$, the first item is still the item that gets split in the new profile. Using the identities $a_1 + a_2 = b_1 + b_2 = 1$ and the assumption that (\mathbf{x}, \mathbf{y}) is a pure Nash equilibrium, we have that

$$\begin{cases} a_1 \left(1 - \frac{1}{2x} - \frac{1}{2x-\delta}\right) + a_2 \leq 0 \implies \delta \geq \frac{4x(x-a_1)}{2x-a_1} \\ b_1 \left(1 - \frac{1}{2x} - \frac{1}{2x+\delta}\right) + b_2 \geq 0 \implies \delta \geq \frac{4x(b_1-x)}{2x-b_1} \end{cases}$$

and we obtain a contradiction.

Case 4: ($x = y = 1/2$). Alice and Bob get allocations $(1, 0)$ and $(0, 1)$, respectively. Let $0 < \delta < \frac{(b_1-b_2)}{b_2}$ and consider the strategy $\mathbf{y}' = (x + \delta, 1 - x - \delta)$ of Bob. Using \mathbf{y}' , Bob gets the allocation $(\frac{1}{\delta+1}, 0)$, which is better than $(0, 1)$. Since $b_1 > b_2$, such δ exists.

As none of the cases 1 – 4 are stable, the procedure has no pure Nash equilibrium. \square

However, we show that Adjusted Winner admits approximate Nash equilibria.

Theorem 5 *Each instance of Adjusted Winner has an ϵ -Nash equilibrium, for every $\epsilon > 0$.*

Proof: Let (\mathbf{a}, \mathbf{b}) be any instance. We show there exists an ϵ -Nash equilibrium in which Alice plays her true valuations and Bob plays a small perturbation of Alice's valuations. More formally, we show there exist $\epsilon_1, \dots, \epsilon_m$, such that an ϵ -equilibrium is obtained when Alice plays $\mathbf{a} = (a_1, \dots, a_m)$ and Bob plays $\tilde{\mathbf{a}} = (\tilde{a}_1, \dots, \tilde{a}_m)$, where $\tilde{a}_i = a_i + \epsilon_i$ for each item $i \in [m]$ and $\sum_{i=1}^m \epsilon_i = 0$. The theorem will follow from the next two lemmas. \square

Lemma 2 *For any pair of strategies $(\mathbf{a}, \tilde{\mathbf{a}})$, where $|a_i - \tilde{a}_i| < \epsilon/m$ for all $i \in [m]$, Alice's strategy is an ϵ -best response.*

Proof: Since the procedure is envy-free, Alice gets at least half of the total value by being truthful regardless of Bob's strategy, so:

$$u_{\mathbf{a}}(AW_A(\mathbf{a}, \tilde{\mathbf{a}})) \geq 1/2.$$

Assume for contradiction that there is another strategy \mathbf{a}' that is a better response for Alice. Then it must hold that

$$u_{\mathbf{a}}(AW_A(\mathbf{a}', \tilde{\mathbf{a}})) > 1/2 + \epsilon.$$

Now since strategies \mathbf{a} and $\tilde{\mathbf{a}}$ are ϵ -close, then $\sum_i |a_i - \tilde{a}_i| < \epsilon$, so it holds that:

$$\begin{aligned} u_{\tilde{\mathbf{a}}}(AW_B(\mathbf{a}', \tilde{\mathbf{a}})) &\leq u_{\mathbf{a}}(AW_B(\mathbf{a}', \tilde{\mathbf{a}})) + \epsilon \\ &= 1 - u_{\mathbf{a}}(AW_A(\mathbf{a}', \tilde{\mathbf{a}})) + \epsilon < 1/2 \end{aligned}$$

But since the allocation must also be envy-free according to Bob's declared valuation profile $\tilde{\mathbf{a}}$, we have that:

$$u_{\tilde{\mathbf{a}}}(AW_B(\mathbf{a}', \tilde{\mathbf{a}})) \geq 1/2$$

The last inequality gives a contradiction. Thus when Bob's strategy is ϵ -close to Alice's truthful strategy \mathbf{a} , Alice's truthful strategy \mathbf{a} is a ϵ -best response, which completes the proof of the lemma. \square

Lemma 3 When Alice plays \mathbf{a} , Bob has an ϵ -best response that is ϵ -close to Alice's strategy.

Proof: Let $\pi = (\pi_1, \dots, \pi_m)$ be a fixed permutation of the items. Then there exist uniquely defined index $l \in \{1, \dots, m\}$ and $\lambda \in [0, 1)$ such that

$$a_{\pi_1} + \dots + a_{\pi_{l-1}} + \lambda a_{\pi_l} = \frac{1}{2} = (1-\lambda)a_{\pi_l} + a_{\pi_{l+1}} + \dots + a_{\pi_m} \quad (1)$$

Note that Adjusted Winner uses lexicographic tie breaking to sort the items when there exist equal ratios $x_i/y_i = x_j/y_j$, for some $i \neq j$. Thus the order π may never appear in an outcome of the procedure when the agents use the same strategies.

However, we show that Bob can approximate the outcome of Equation (1) arbitrarily well. We have two cases:

Case 1: $\lambda \in (0, 1)$. Then there exist $\epsilon_1, \dots, \epsilon_m$ such that the following conditions hold:

- $|\epsilon_j| < \min\left(\frac{\epsilon}{m}, \frac{2\lambda a_{\pi_l}}{m}\right)$, for all $j \in [m]$,
- the items are strictly ordered by π : $\frac{a_{\pi_1}}{a_{\pi_1} + \epsilon_{\pi_1}} > \dots > \frac{a_{\pi_m}}{a_{\pi_m} + \epsilon_{\pi_m}}$,
- $\sum_{j=1}^m \epsilon_j = 0$,
- it's still item π_l that gets split, in a fraction $\delta \in (0, 1)$ close to λ ; that is, $|\lambda - \delta| < \frac{\epsilon}{b_{\pi_l}}$.

Informally, Bob plays a perturbation of Alice's truthful strategy inducing ordering π on the items (with no ties) and splits item π_l in a fraction close to λ .

Case 2: $\lambda = 0$. Again, there are $\epsilon_1, \dots, \epsilon_m$, such that the following conditions hold:

- $\epsilon_j < \min\left(\frac{\epsilon}{m}, \frac{a_{\pi_l}}{m}\right)$ for all $j \in [m]$,
- the item order is π : $\frac{a_{\pi_1}}{a_{\pi_1} + \epsilon_{\pi_1}} > \dots > \frac{a_{\pi_m}}{a_{\pi_m} + \epsilon_{\pi_m}}$,
- $\sum_{j=1}^m \epsilon_j = 0$,
- item π_l is split in a ratio δ close to zero: $|\delta| < \frac{\epsilon}{b_{\pi_l}}$.

Thus Bob can approximate the outcome of Equation (1).

Now consider any ϵ -best response \mathbf{y} of Bob; this induces some permutation of the items according to the ratios. If \mathbf{y} is ϵ -close to the strategy of Alice we are done. Otherwise, Bob could change his strategy to be ϵ -close to the strategy of Alice while inducing the same permutation. This will only improve his utility as the boundary line moves to the left. \square

It can be observed that there is at least one other ϵ -Nash equilibrium, at strategies $(\mathbf{b}, \tilde{\mathbf{b}})$, where $\tilde{\mathbf{b}}$ is a perturbation of Bob's truthful profile.

Discrete Strategies

Even though the continuous procedure is not guaranteed to have pure Nash equilibria, this does not imply that the discrete variant should also fail to have pure Nash equilibria. However we do find that this is indeed the case.

Theorem 6 Adjusted Winner with discrete strategies is not guaranteed to have pure Nash equilibria.

Proof: Consider a game with 4 items and 7 points, where Alice and Bob have valuations $(1, 1, 2, 3)$ and $(2, 3, 1, 1)$, respectively. This game does not admit a pure Nash equilibrium; this fact can be verified with a program that checks all possible configurations. \square

Our next theorem shows that an ϵ -Nash equilibrium always exists in the discrete case if the number of points is set adequately, such that the agents can approximately represent their true valuations.

Theorem 7 For any profile (\mathbf{a}, \mathbf{b}) and any $\epsilon > 0$, there exists P' such that the procedure has an ϵ -Nash equilibrium when the agents are given P' points.

Proof: Let $\epsilon > 0$, and consider any profile (\mathbf{a}, \mathbf{b}) with denominator P . Then if we interpret (\mathbf{a}, \mathbf{b}) as a profile for the continuous setting, we get a $\epsilon/2$ -Nash equilibrium $(\mathbf{a}, \tilde{\mathbf{a}})$ from Theorem 5, where $\tilde{a}_j = a_j + \epsilon_j$, for all $j \in [m]$.

Recall that $a_j, b_j \in \mathbb{Q}$; where $a_j = \frac{s_j}{P}$ and $b_j = \frac{t_j}{P}$, for some $s_j, t_j \in \mathbb{N}$. We can find a rational number $\epsilon'_j = \frac{q_j}{r_j}$ (with $q_j, r_j \in \mathbb{N}$) that approximates ϵ_j within $\frac{\epsilon}{2m}$ for each $j \in [m]$, and such that the ordering of the items induced by the ratios $\frac{a_j}{a_j + \epsilon_j}$ is the same as the one given by $\frac{a_j}{a_j + \epsilon'_j}$. Define $\tilde{\mathbf{a}}'$ such that $\tilde{a}'_j = a_j + \epsilon'_j$.

It follows that $(\mathbf{a}, \tilde{\mathbf{a}}')$ is an ϵ -Nash equilibrium with $a_j, \tilde{a}'_j \in \mathbb{Q}$, for all $j \in [m]$. Thus whenever the agents have a denominator of $P' = P \cdot \prod_{j=1}^m r_j$, the strategy profiles $(\mathbf{a}, \tilde{\mathbf{a}}')$ can be represented in the discrete procedure, so by giving P' points to the agents, there exists an ϵ -Nash equilibrium. \square

4.2 Informed Tie-Breaking

If the tie-breaking rule is not independent of the valuations, then both the discrete and continuous variants of Adjusted Winner have exact pure Nash equilibria. The deterministic tie-breaking rule under which this is possible is the one in which one of the agents, for example Bob, is allowed to resolve ties by sorting them in the best possible order for him. Bob can compute the optimal order as outlined in the next definition.

Definition 1 (Informed Tie-Breaking) Let there be a fixed agent, for example Bob. Given any strategies (\mathbf{x}, \mathbf{y}) , for each permutation π , let $l_\pi \in [m]$ and $\lambda_\pi \in [0, 1)$ be the uniquely defined item and fraction for which:

$$x_{\pi_1} + \dots + x_{\pi_{l-1}} + \lambda x_{\pi_l} = (1-\lambda)y_{\pi_l} + y_{\pi_{l+1}} + \dots + y_{\pi_m}$$

Let π^* be an optimal permutation with respect to (\mathbf{x}, \mathbf{y}) , namely $\pi^* \in \arg \max_{\pi} (1-\lambda)y_{\pi_l} + y_{\pi_{l+1}} + \dots + y_{\pi_m}$. Then under informed tie-breaking, the procedure resolves ties in the order given by π^* .

Note that there might be more than one choice of π^* and Bob picks any fixed one. Now we can state the equilibrium existence theorems.

Theorem 8 Adjusted Winner with continuous strategies and informed tie-breaking is guaranteed to have a pure Nash equilibrium.

Proof: We show that the profile (\mathbf{a}, \mathbf{a}) is an exact equilibrium. By envy-freeness of the procedure, Alice gets at least half of the points at this strategy profile. Moreover, she cannot get strictly above half, since that would violate envy-freeness from the point of view of Bob's declared valuation, which is also \mathbf{a} . Thus Alice's strategy is a best response. As argued in Theorem 5 and 7, there exists an optimal permutation π^* such that by playing \mathbf{a} and sorting the items in the order π^* , Bob can obtain the best possible utility (and as mentioned in Lemma 3, this value is achievable at these strategies). \square

Similarly, it can be shown that the strategy profile (\mathbf{a}, \mathbf{a}) is a pure Nash equilibrium in the discrete procedure.

Theorem 9 *Adjusted Winner with discrete strategies and informed tie-breaking is guaranteed to have a pure Nash equilibrium.*

Proof: Consider the strategy profile (\mathbf{a}, \mathbf{a}) . From Theorem 8, this is a Nash equilibrium in the continuous case. Since the strategy space in the discrete procedure is more restricted, there are no improving deviations here either, and so the theorem follows. \square

5 Efficiency and Fairness of Equilibria

Having examined the existence of pure Nash equilibria in Adjusted Winner, we now study the fairness and efficiency of exact equilibria. For fairness, we observe that following. The reason is that by reporting truthfully, each agent guarantees at least $1/2$ utility [Barbanel and Brams, 2014].

Theorem 10 *All the pure Nash equilibria of Adjusted Winner are envy-free with respect to the true valuations of the agents.*

For efficiency, we use the well known measure of the *Price of Anarchy* [Koutsoupias and Papadimitriou, 1999; Nisan *et al.*, editors 2007].

First, the *social welfare* of an allocation W is defined as the sum of the agents' utilities: $SW(W) = u_A(W_A) + u_B(W_B)$. Then the Price of Anarchy is defined as the ratio between the maximum social welfare and the welfare of the worst-case pure Nash equilibrium.

Our main findings are that when the procedure is equipped with an informed tie-breaking rule (i) all the pure Nash equilibria are Pareto optimal with respect to the true valuations and (ii) the price of anarchy is constant; that is, each pure Nash equilibrium achieves at least 75% of the optimal social welfare.

From the previous discussion, it can be observed that in every exact equilibrium, Alice and Bob copy each other's strategy.

Theorem 11 *Every Nash equilibrium of Adjusted Winner occurs at "symmetric" strategy profiles, of the form (\mathbf{x}, \mathbf{x}) .*

We start with a lemma.

Lemma 4 *Let (\mathbf{x}, \mathbf{x}) be a pure Nash equilibrium of Adjusted Winner with informed tie-breaking and let π^* be the permutation that Bob chooses. Then, among all possible permutations, π^* maximizes Alice's utility.*

Proof: Assume by contradiction that there exists a permutation π that gives Alice a strictly larger utility; let α be her marginal increase from π^* to π . As discussed in Section 4, Alice can find appropriate constants $\epsilon_1, \dots, \epsilon_m$ such that $AW(\mathbf{x}', \mathbf{x})$ with $\mathbf{x}' = (x_1 + \epsilon_1, \dots, x_m + \epsilon_m)$ orders the items by π and the allocations $AW(\mathbf{x}, \mathbf{x})$ and $AW(\mathbf{x}', \mathbf{x})$ differ only in the allocation of the split item by δ . Moreover, by continuity of the strategies, for each α , there exist ϵ_i 's such that δ is small enough for $AW(\mathbf{x}', \mathbf{x})$ to be better for Alice than $AW(\mathbf{x}, \mathbf{x})$. \square

Next we show that all equilibria are Pareto optimal.

Theorem 12 *All the pure Nash equilibria of Adjusted Winner with informed tie-breaking are Pareto optimal with respect to the true valuations \mathbf{a} and \mathbf{b} .*

Proof: Let (\mathbf{x}, \mathbf{x}) be a pure Nash equilibrium of Adjusted Winner under informed tie-breaking and let l be the item that gets split (if any, otherwise the item to the left of the boundary line). Order Alice's items decreasing order of ratios a_i/x_i and Bob's items in increasing order of ratios b_i/x_i . Since (\mathbf{x}, \mathbf{x}) is a pure Nash equilibrium, by Lemma 4, both agents are getting their maximum utility over all possible tie-breaking orderings of items. This means that for every item $i \leq l$ and every item $j \geq l$ with $i \neq j$, it holds that

$$\frac{a_j}{x_j} \geq \frac{a_i}{x_i} \quad \text{and} \quad \frac{b_i}{x_i} \geq \frac{b_j}{x_j} \Rightarrow \frac{a_i}{x_i} \cdot \frac{b_j}{x_j} \leq \frac{a_j}{x_j} \cdot \frac{b_i}{x_i}$$

which by Lemma 1, implies that $AW(\mathbf{x}, \mathbf{x})$ is Pareto optimal. \square

The Pareto optimality of a strategy profile has a direct implication on the social welfare achieved at that profile.

Theorem 13 *The Price of Anarchy of Adjusted Winner is $4/3$.*

Proof: Let (\mathbf{x}, \mathbf{y}) be any pure Nash equilibrium and let OPT_A and OPT_B be the utilities of Alice and Bob respectively in the optimal allocation. Since $AW(\mathbf{x}, \mathbf{y})$ is Pareto optimal by Theorem 12, the allocation for at least one of the agents, (e.g. Alice), is at least as good as that of the optimal allocation. In other words, $u_A(AW(\mathbf{x}, \mathbf{y})) \geq OPT_A$. On the other hand, since $AW(\mathbf{x}, \mathbf{y})$ is envy-free, Bob's utility from $AW(\mathbf{x}, \mathbf{y})$ is at least $1/2$ which is at least $\frac{1}{2}OPT_B$. Overall, the social welfare of $AW(\mathbf{x}, \mathbf{y})$ is at least $OPT_A + \frac{1}{2}OPT_B$ and the ratio is minimized when OPT_A and OPT_B are minimum. Since $OPT_A \geq OPT_B \geq 1/2$, the ratio is at least $4/3$.

The bound is (almost) tight, given by the following simple instance with two items. Let $\mathbf{a} = (1 - \epsilon, \epsilon)$ and $\mathbf{b} = (\epsilon, 1 - \epsilon)$ and consider the strategy profile $\mathbf{x} = (\epsilon, 1 - \epsilon)$ and $\mathbf{y} = (\epsilon, 1 - \epsilon)$. It is not hard to see that \mathbf{x}, \mathbf{y} is a pure Nash equilibrium for Alice breaking ties. The social welfare of the optimal allocation is $2 - 2\epsilon$. In the allocation of Adjusted Winner, Alice wins the first item and the second item is split (almost) in half. The social welfare of the mechanism is $1 + \frac{1}{2} + o(\epsilon)$ and the approximation ratio is (almost) $4/3$. As ϵ grows smaller, the ratio becomes closer to $4/3$. \square

6 Future Work

According to Foley [1967], the quintessential characteristics of fairness are envy-freeness and Pareto optimality. We show that Adjusted Winner is guaranteed to have pure Nash equilibria, which satisfy both of these fairness notions. This attests to the usefulness and theoretical robustness of the procedure. A very interesting direction for future work is to study the *imperfect information* setting, as the Nash equilibria studied here require the agents to have full information of each other's preferences.

Acknowledgments

NICTA is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program.

Simina Brânzei, Aris Filos-Ratsikas and Søren Kristoffer Stiil Frederiksen were supported by the Sino-Danish Center for the Theory of Interactive Computation, funded by the Danish National Research Foundation and the National Science Foundation of China (under the grant 61061130540), and by the Center for research in the Foundations of Electronic Markets (CFEM), supported by the Danish Strategic Research Council. Simina Brânzei acknowledges additional support from an IBM Ph.D. fellowship.

The authors thank Steven Brams for useful pointers to the literature and Edith Elkind for helpful technical discussion.

References

- J. B. Barbanel and S. J. Brams. Two-person cake cutting: The optimal number of cuts. *The Mathematical Intelligencer*, 36(3):23–35, 2014.
- S. J. Brams and A. D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.
- S. J. Brams and A. D. Taylor. A procedure for divorce settlements. *Issue Mediation Quarterly Mediation Quarterly*, 13(3):191–205, 1996.
- S. J. Brams and A. D. Taylor. *The Win-Win Solution: Guaranteeing Fair Shares to Everybody*. Norton, 2000.
- S. J. Brams and J. M. Togman. Camp david: Was the agreement fair? *Conflict Management and Peace Science*, 15(1):99–112, 1996.
- S. J. Brams, M. Feldman, J. Morgenstern, J. K. Lai, and A. D. Procaccia. On maxsum fair cake divisions. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, pages 1285–1291. AAAI Press, 2012.
- Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J. A. Rodríguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica*, 30:3–31, 2006.
- M. Dall'Aglio and R. Mosca. How to allocate hard candies fairly. *Mathematical Social Sciences*, 54(3):218–237, 2007.
- T. E. Daniel and J. E. Parco. Fair, efficient and envy-free bargaining: An experimental test of the brams-taylor adjusted winner mechanism. *Group Decision and Negotiation*, 14(3):241–264, 2005.
- D. B. H. Denoon and S. J. Brams. Fair division: A new approach to the spratly islands controversy. *International Relations*, 2(2):303–329, 1997.
- D. K. Foley. Resource allocation and the public sector. *Yale Econ Essays, Vol 7, No 1, pp 45-98, Spring 1967. 7 Fig, 13 Ref.*, 1967.
- J. Karp, A. M. Kazachkov, and A. D. Procaccia. Envy-free division of sellable goods. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, pages 728–734. AAAI Press, 2014.
- E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *STACS 99*, pages 404–413. Springer, 1999.
- J. R. Lax. Fair division: A format for the debate on the format of debates. *Political Science and Politics*, 32(1):45–52, 1999.
- G. M. Levy. Resolving real estate disputes. *Real Estate Issues*, 1999.
- K. Madani. Game theory and water resources. *Journal of Hydrology*, 381:225–238, 2010.
- T. G. Massoud. Fair division, adjusted winner procedure (aw), and the israeli-palestinian conflict. *Journal of Conflict Resolution*, 44(3):333–358, 2000.
- H. Moulin. *Fair Division and Collective Welfare*. MIT Press, 2003.
- B. Nalebuff. Review of the win-win solution: Guaranteeing fair shares to everybody. *Journal of Economic Literature*, 39:125–127, 2001.
- N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, (editors) 2007.
- A. D. Procaccia. Cake cutting: Not just child's play. *Communications of the ACM*, 56(7):78–87, 2013.
- M. G. Raith. Fair-negotiation procedures. *Mathematical Social Sciences*, 39:303–322, 2000.
- J. M. Robertson and W. A. Webb. *Cake Cutting Algorithms: Be Fair If You Can*. A. K. Peters, 1998.
- A. D. Taylor and A. M. Pacelli. *Mathematics and Politics*. Springer, 2008.
- P. H. Young. *Equity in theory and practice*. Princeton University Press, 1994.