# Long-term Fairness with Bounded Worst-case Losses 

Gabriel Balan and Dana Richards and Sean Luke<br>Department of Computer Science, George Mason University<br>4400 University Drive, MSN 4A5, Fairfax, VA, 20030, USA<br>\{gbalan, richards, sean\} @cs.gmu.edu


#### Abstract

How does one repeatedly choose actions so as to be fairest to the multiple beneficiaries of those actions? We examine approaches to discovering sequences of actions for which the worst-off beneficiaries are treated maximally well, then secondarily the second-worst-off, and so on. We formulate the problem for the situation where the sequence of action choices continues forever; this problem may be reduced to a set of linear programs. We then extend the problem to situations where the game ends at some unknown finite time in the future. We demonstrate that an optimal solution is NP-hard, and present two good approximation algorithms.


## Introduction

Consider the problem of repeatedly assigning two AI professors to teach two classes offered by their department each semester. One class is much harder than the other one, so during any single semester any one-to-one assignment is unfair to one of the professors. One fair solution would be for the professors to teach both classes together. But assuming this requires more overall effort than teaching the classes separately, this solution would be inefficient over the long run.

There is of course a better solution. If the two professors instead took turns teaching the hard class, then in the long run their average utilities would be more fair than in the one-to-one assignments and more efficient than in the sharing assignment. This is the rough idea behind long-term fairness: repeated interactions offer opportunities for improved efficiency and fairness over the single interaction scenario. But in general the solutions will not be as simple as alternating assignments (e.g. suppose we extended the previous example with multiple assignments, involving many professors and classes).

The research presented here examines the following framework: there are a number of beneficiaries (e.g. professors), which receive different rewards from each of a finite set of actions (e.g. class assignments). The actions are chosen with replacement, and actions chosen early do not restrict what actions can be chosen later, or their rewards. This framework, borrowed from (Verbeeck, Parent,

[^0]and Nowé 2003), is very similar to the repeated normal-form game framework from game theory, except there is a single decision maker that chooses actions for the good of all beneficiaries.
We are ultimately interested in the game-theoretic multiple decision maker case. But it turns out that the single decision maker situation has been inadequately studied, and so our focus here is on that case.

We define a beneficiary's utility as the average of all rewards he received in the past. We will use the term utility profile to refer to the vector of utilities (one utility per beneficiary) that they derive from past actions.

In order to compare utility profiles, we use a fairness concept called leximin. Leximin is a total-order relation defined as follows: a utility profile $U_{1}$ is preferable to another utility profile $U_{2}$ if the beneficiary with the minimum utility in $U_{1}$ is better-off than the minimum utility beneficiary in $U_{2}$; ties are broken by comparing the utilities of the the secondworst beneficiary in each utility profile, then third-worst, and so on. From a procedural point of view, one should sort in increasing order the two utility profiles, and compare the sorted versions lexicographically. We choose leximin for our work because it is widely used in the literature, it is well understood, and it incorporates a certain measure of efficiency (a leximin-optimal profile is also Pareto-optimal), making it a good fit for long-term fairness.

In this paper we derive upper bounds for this notion of utility in infinitely repeated games and propose a particular approach to approximation solutions for more realistic cases (i.e. a finite sequence of actions).

## Infinite-length Games

An action may give high rewards to some beneficiaries and low rewards to others, so sticking to just one action might be unfair to some beneficiaries, and thus leximin-undesirable. However, if beneficiaries receive different rewards from different actions we may be able to improve all beneficiaries’ reward averages by playing a combination of actions.

One approach to playing combinations is to use a periodic, repeated sequence of actions. In previous literature (Verbeeck, Parent, and Nowé 2003), distributed algorithms for discovering such sequences found suboptimal ones. We will show optimal solutions, albeit with non-distributed algorithms.

Periodic sequences are intuitively appealing, but even if one can find the periodic sequence with the best limit, that can still be suboptimal. We will show at the end of this section that in some problem instances with irrational coefficients there might exist infinite non-periodic sequences that achieve, at the limit, leximin-superior utility profiles to any utility profile achievable by a periodic sequence. Our algorithms are guaranteed to produce sequences converging to the optimal utility profile, and, whenever possible, those sequences are periodic.

In this section we will (1) prove that convergent sequences (periodic or non-periodic) are sufficient and (2) identify the class of sequences with the leximin-optimal limit-point. In the following sections we will propose additional requirements to impose on this class of sequences and then provide algorithms that produce sequences satisfying these requirements.
Problem 1. Base Problem Let there be a set $\mathbb{A}$ of $n_{a}$ actions affecting a set $\mathbb{B}$ of $n_{b}$ beneficiaries through the reward function $R: \mathbb{A} \times \mathbb{B} \rightarrow \mathbb{R}$. Let $\mathbb{S}$ be the set of infinite sequences of actions from $\mathbb{A}$ :

$$
\mathbb{S}=\left\{S=\left\langle s_{1}, s_{2}, \ldots\right\rangle \mid \forall i \in \mathbb{N}: s_{i} \in \mathbb{A}\right\} .
$$

Let $\mathbb{U}$ be the set of all possible utility profiles (vectors) achievable from following any sequence for any number of time steps:

$$
\begin{aligned}
& \text { ne steps: } \\
& \qquad \mathbb{U}=\left\{U \in \mathbb{R}^{n_{b}} \mid \exists t \in \mathbb{N}, \exists S \in \mathbb{S}: U_{b}=\frac{1}{t} \sum_{j=1}^{t} R\left(s_{j}, b\right)\right\} .
\end{aligned}
$$

The goal of the problem is to find $U^{\star}=\sup (\mathbb{U})$, the supremum (least upper bound) with respect to leximin over the set of all achievable utility profiles.

In order to solve this problem we focus on a subset of "well-behaved" sequences that achieve the entire set of possible utility profiles $\mathbb{U}$. We then provide a mapping of that subset into the $n_{a}$-dimensional simplex which can be searched efficiently using linear-programming-based algorithms from the literature.

To make things easier, we refer to the elements of $\mathbb{A}$ as $\left\{1 \ldots n_{a}\right\}$ and the elements of $\mathbb{B}$ as $\left\{1 \ldots n_{b}\right\}$. Let $\mathbb{S}^{\prime} \subset \mathbb{S}$ be the set of all sequences $S$ where the proportions of the $n_{a}$ actions converge. Formally:

$$
\mathbb{S}^{\prime}=\left\{S \in \mathbb{S} \mid \forall j \in \mathbb{A}: \exists \lim _{t \rightarrow \infty} \frac{1}{t} k_{j}\left(S_{t}\right)\right\}
$$

where $S_{t}$ is the subsequence of $S$ consisting of the first $t$ elements and $k$ is the count function (so $k_{j}\left(S_{t}\right)$ is equal to the number of times action $j$ is used in the first $t$ positions of sequence $S$ ). We denote with $F(S)$ the vector of action proportions (or fractions) for $S$ (i.e. $F_{j}(S)=\lim _{t \rightarrow \infty} \frac{1}{t} k_{j}\left(S_{t}\right)$ ). Let $U\left(S_{t}\right)$ be the vector of utilities achieved after following the first $t$ steps of sequence $S$ (i.e. $U$ 's component for beneficiary $b$ is $U_{b}\left(S_{t}\right)=\frac{1}{t} \sum_{i=1}^{t} R\left(s_{j}, b\right)$ ). Because the action proportions converge, the sequence of utility vectors $\left\langle U\left(S_{1}\right), U\left(S_{2}\right), \ldots\right\rangle$ also converges component by component; we denote its limit point by the vector $U(S)$ :

$$
\begin{equation*}
U_{b}(S)=\lim _{t \rightarrow \infty} U_{b}\left(S_{t}\right)=\sum_{j=1}^{n_{a}} F_{j}(S) R\left(s_{j}, b\right) \tag{1}
\end{equation*}
$$

It can be shown that $\mathbb{U}^{\prime}$, the set of utility profiles achievable by sequences in $\mathbb{S}^{\prime}$, is equal to $\mathbb{U} .^{\ddagger}$ Moreover, if $\mathbb{U}^{\prime \prime}$ is the set of limit points of $\mathbb{U}^{\prime}\left(\mathbb{U}^{\prime \prime}=\left\{U \mid \exists S \in \mathbb{S}^{\prime}: U=U(S)\right\}\right)$, then we can prove that: ${ }^{\ddagger}$

$$
\begin{equation*}
\sup \left(\mathbb{U}=\mathbb{U}^{\prime}\right)=\sup \left(\mathbb{U}^{\prime \prime}\right) \tag{2}
\end{equation*}
$$

Because $U(S)$ depends only on $F(S)$, there is an obvious one-to-one correspondence between $\mathbb{U}^{\prime \prime}$ and the $n_{a^{-}}$ dimensional unit simplex.

Once reformulated as an optimization problem over a compact and convex set, the problem becomes readily solvable with the algorithm proposed in (Potters and Tijs 1992). In our particular case, the algorithm consists of solving $O\left(n_{a}^{2}\right)$ linear programs (LPs), but approximate results based on a floating-point fixed-length representation might be required to make sure the complexity of this algorithm does not dominate that of the algorithms we propose here. The algorithm produces $U^{\star}$, the unique leximin-optimal utility vector, and $F^{\star}$, a point (not necessarily unique) inside the simplex such that: $\forall b \in \mathbb{B}: \sum_{i=1}^{n_{a}} R(b, i) \times F_{i}^{\star}=U_{b}^{\star}$. Since $U^{\star}$ is the leximin-maximal utility vector in $\mathbb{U}^{\prime \prime}$ (by the algorithm's guarantee), and it is unique, it implies that $U^{\star}=\sup \left(\mathbb{U}^{\prime \prime}\right)$, and, by Equation 2:

$$
\begin{equation*}
U^{\star}=\sup (\mathbb{U}) \tag{3}
\end{equation*}
$$

We revisit the existence of a periodic sequence for a given optimal utility profile solution $U^{\star}$. If such a periodic policy exists, let $c_{j}$ denote the number of times of action $j$ appears in one period (we call $c_{j}$ the multiplicity of action $j$ ). Then by normalizing the vector $c=\left[c_{1} \ldots c_{n_{a}}\right]$ one should obtain a valid $F^{\star}$ vector. If all rewards are rational then there must exist an $F^{\star}$ with rational components, and hence a periodic sequence. ${ }^{\ddagger}$ However, irrational rewards could mean there is no $F^{\star}$ with rational coefficients, in which case there is no optimal periodic sequence.

For example, if both professors get a reward of 0 from teaching the hard class, but teaching the easy class gives one professor a reward of 1 and the other a reward of $\sqrt{2}$, then $U^{\star}=\left[\frac{\sqrt{2}}{1+\sqrt{2}}, \frac{\sqrt{2}}{1+\sqrt{2}}\right]$ and there is a unique $F^{\star}=\left[\frac{\sqrt{2}}{1+\sqrt{2}}, \frac{1}{1+\sqrt{2}}\right]$. We make the observation that depending on the values in $F^{\star}$, the multiplicities could be arbitrarily large. Thus, one can see the non-periodic sequences due to $F^{\star}$ having (some) irrational components as special cases of periodic sequences where the period length is infinite.

## Finite-length Games

Solving the Base Problem provides $U^{\star}$, the least upper bound over the set of achievable utility profiles. If the game is guaranteed to last forever, one could be satisfied with the goal of finding a sequence that achieves $U^{\star}$ at the limit. But in more practical applications, one must consider the implications of having the process end after a finite number of steps, or more generally, having the length of the sequence of actions drawn from some probability distribution.

[^1]

Figure 1: Risks as functions of time for the two beneficiaries in Example 1 (the solid line for the first and the dotted line for the second) as produced by various periodic sequences.

Consider the professor-assigment example from the introduction: the first actions corresponds to the assignment where the first professor teaches the easy class, the second action corresponds to the assignment where the second professor teaches the easy class, and the third action is the assignment where they teach the classes together:

Example 1.


The utility least upper bound for Example 1 is $U^{\star}=[5,5]$, which can only be achieved through $F^{\star}=[0.5,0.5,0]$, and which coincides with the optimal solution if the game lasts for an even number of steps: use actions 1 and 2 in equal proportions. However, one should use action 3 if the game lasts for only one round. This shows that the optimal sequence of decisions depends on the duration of the game, and so if the duration is not known in advance, some sort of tradeoff might be required.

One could choose an action stochastically, using the values in $F^{\star}$ as probabilities. This approach produces the leximin-optimal expected rewards (equal to $U^{\star}$ ). This sort of guarantee is usually referred to in the economics literature as ex-ante fairness (ex-ante means "beforehand" in latin). If all beneficiaries are risk-neutral, a mechanism choosing repeatedly from the $F^{\star}$ distribution promises leximin-optimal expected utilities before the process starts, but not necessarily leximin-optimal actual utilities when the game ends. For Example 1 this approach would flip a coin between actions 1 and 2; although the expected rewards (and utilities) are equal to $U^{\star}$, if the game ends after two rounds, there is a $50 \%$ chance that one beneficiary accumulated 10 reward units and the other none. Moreover, if the game continues, the probability that the difference between the two will shrink is equal to the probability that the difference will grow. Thus, the weakness of this ex-ante approach would be its lack of concern with paying reparations for unfairness resulting from past randomness.

One can address the weaknesses of the previous method with a deterministic approach, aiming to find a sequence which leximin-optimizes beneficiaries' expected utilities weighted by the probabilities in the distribution of game
lengths. Unlike the previous method, this could use action 3 in Example 1, provided there is a high enough probability that the game will end after one step.

Both previous methods assume risk-neutral beneficiaries. Given that leximin is a risk-averse fairness concept ("no one is left behind"), it makes sense to instead focus on a riskaverse solution approach: minimize the largest amount a beneficiary risks losing due to the game ending prematurely. We define the risk of beneficiary $k$ at time step $t$ while executing sequence $S$ as the difference between the rewards accumulated by beneficiary $k$ during the first $t$ steps, and the amount he was entitled to, which is $t \times U_{k}^{\star}$.

$$
\begin{equation*}
\operatorname{Risk}_{b}(S, t)=\sum_{i=1}^{t} R\left(s_{i}, b\right)-t \times U_{b}^{\star} \tag{4}
\end{equation*}
$$

We use the term worst risk (WR) of a sequence $S$ to mean a lower bound on all risk values, regardless of beneficiary (i.e. $W R \leq \operatorname{Risk}_{b}(S, t), \forall t \in \mathbb{N}$ and $\forall b \in \mathbb{B}$ ). The worst risk is a negative value, and its absolute value is an upper bound on the largest amount that any beneficiary can lose.

This approach is not meant to replace expectation optimization; we are simply adding another tool to the arsenal available for this problem. It is reasonable to expect the computational complexity of the expectation-optimization approach to grow with the maximum game length, and so our proposed solution is particularly suitable to scenarios with very large game durations (our approach is insensitive to the distribution of game durations). Furthermore, lowerbounding the worst risk will always make the utilities converge to $U^{\star}$ as $t \rightarrow \infty^{\star}$

For convenience, we define for each action $j$ a vector $X_{j}=$ [ $X_{j, 1} \ldots X_{j, n_{b}}$ ], where $X_{j, b}=R(j, b)-U_{b}^{\star}$, the amount action $j$ changes the risk of beneficiary $b$.

$$
\begin{gather*}
\operatorname{Risk}_{b}(S, t)=\sum_{i=1}^{t}\left[R\left(s_{i}, b\right)-U_{b}^{\star}\right]=\sum_{i=1}^{t} X_{s_{i}, b}  \tag{5}\\
\forall b \in \mathbb{B}: \sum_{j=1}^{n_{a}} X_{j, b} \times F_{j}^{\star}=0 \tag{6}
\end{gather*}
$$

Let us illustrate the concept of risk using Example 1. The $X$ vectors are as follows: $X_{1}=[5,-5], X_{2}=[-5,5]$, and $X_{3}=[-4,-4]$. We will compare the following periodic ac-
tion sequences (see Figure 1):

$$
\begin{aligned}
& S_{1}=\langle 1,2,1,2, \ldots\rangle \\
& S_{2}=\langle 2,1,2,1, \ldots\rangle \\
& S_{3}=\langle 1,1,2,2,1,1,2,2, \ldots\rangle \\
& S_{4}=\langle 1,2,2,1,1,2,2,1, \ldots\rangle
\end{aligned}
$$

Sequence $S_{1}$ makes the first beneficiary's risks equal to 5 on odd steps and 0 on even steps, while the second beneficiary's risks are 0 on odd steps and -5 on even ones. Therefore the second beneficiary risks coming up 5 units short if the game stops after an odd number of steps and breaks even otherwise. Sequence $S_{2}$ has identical effects as $S_{1}$, but to different beneficiaries, so the two are equally good. Using sequence $S_{3}$, the second beneficiary risks losing as much as 10 units, so we prefer $S_{1}$ (or $S_{2}$ ) to $S_{3}$.

Note that using sequence $S_{1}$ makes the second beneficiary take all the "bad" (negative) risks, while the first beneficiary takes only "good" (positive) risks. Although sequence $S_{1}$ is as fair as possible with respect to average utilities, one cannot help but notice a "second degree" unfairness. Instead, one might consider the goal of optimizing cumulative risks (e.g. sequence $S_{4}$ alternates which beneficiary is taking negative risks). This leads to the following paradox: although intuitively one prefers $S_{4}$ to $S_{1}$ or $S_{2}$, the worst risks for all beneficiaries during sequences $S_{1}$ or $S_{2}$ (i.e. -5 and 0 ) are leximin superior to the worst risks during $S_{4}$ (i.e. -5 and $-5)$. This suggests that one might consider optimizing only the "min" of the worst risks, instead of leximin optimizing all beneficiaries' worst risks.

Problem 2. Worst Risk Maximization Problem Given the setup in the Base Problem, find an infinite sequence of actions $S$ with a maximal worst risk.
Theorem 1. Problem 2 is intractable ${ }^{\ddagger}$ (since even finding $a$ sequence of a given finite length is NP-hard).

## Algorithms

The most intuitive solution to optimizing worst risk(s) is to try to keep all risks as large as possible at all times, i.e. greedily choose the action that leximin optimizes risks during the next step (choose the action with the best immediate effects). We will refer to this strategy as GR (Greedy leximin-optimizing next-step Risks). One weakness of this approach is that it assumes the world ends at the next step. In Example 1 it will always choose action 3, leading to arbitrarily bad risks for both beneficiaries if the game lasts long enough. Note that $F_{3}^{\star}=0$, meaning action 3 should not be played at all. One can easily extend the greedy heuristic to ignore the "unusable" actions, but that is not enough to prevent risks from getting arbitrarily bad. Consider Example 2, with $U^{\star}=[240,240,240]$ and unique $F^{\star}=\left[\frac{4}{15}, \frac{6}{15}, \frac{5}{15}\right]$. For this problem GR will choose action 1 repeatedly, leading to arbitrarily bad risks for the first and last beneficiaries. The reason is that action 2 and 3 hurt one of those beneficiaries more than action 1 hurts any of them, and GR lacks the look-ahead to see the benefits of using actions 2 and 3.

| Example 2. |  |  | Beneficiaries |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  | b1 | b2 | b3 |
|  |  | 1 | -20 | 30 | -20 |
|  | Actions | 2 | 60 | 30 | -40 |
|  |  | 3 | -56 | -60 | 64 |

The algorithms we propose in this paper are based on the following observation. All beneficiaries' risks are zero at time $t$ if the number of times each action $j$ was used up to time $t$ are all proportional to the corresponding components of some $F^{\star}$ vector (i.e. $k_{j}\left(S_{t}\right)=F_{j}^{\star} \times t, \forall j$ ). Ideally all $k_{j}\left(S_{t}\right)=F_{j}^{\star} \times t$ at all times, but since the $k_{j}$ functions only take integer values, that is not always possible. It is intuitive that the risks cannot get very bad at any time $t^{\prime}$ if the relative counts $\left(k_{j} / t^{\prime}\right)$ for how much each action was used up to time $t^{\prime}$ are close enough to the corresponding values in $F^{\star}$. There are many ways one can construct sequences $S$ to keep the $k_{j}\left(S_{t}\right)$ values close enough to the $F_{j}^{\star} \times t$ values; we will present two simple methods, and derive risk upper and lower bounds for both of them.

Our methods completely ignore the actions $j$ which are not "used" in $F^{\star}$ (i.e. $F_{j}^{\star}=0$ ). Let $n_{a}^{\star}$ be the number of actions "used" in $F^{\star}$; for simplicity, we rename the actions (reorder the dimensions of the simplex) such that all actions used in $F^{\star}$ come before the other ones: $\forall j \in\left\{1 \ldots n_{a}^{\star}\right\}: F_{j}^{\star}>$ 0 and $\forall j \in\left\{n_{a}^{\star}+1 \ldots n_{a}\right\}: F_{j}^{\star}=0$.

Method 1 Intuitively, this method chooses an action that, up to that point, has been used the least relative to how often the action should have been used. Consequently, an action $j$ can be chosen at time $t+1$ if it has the smallest ratio $\frac{k_{j}\left(D_{t}\right)}{F_{j}^{\star} \times t}$, where $D$ is the sequence of decisions produced by this method (so $D_{t}$ contains all its past decisions). Without affecting the decision process, one can eliminate $t$ from the denominator. We formally describe this method with the following notation:

$$
\begin{equation*}
d_{t} \in\left\{j \leq n_{a}^{\star} \mid \forall i \leq n_{a}^{\star}: \frac{k_{j}\left(D_{t-1}\right)}{F_{j}^{\star}} \leq \frac{k_{i}\left(D_{t-1}\right)}{F_{i}^{\star}}\right\} \tag{7}
\end{equation*}
$$

where all functions $k_{j}$ are extended such that $k_{j}\left(D_{0}\right)=0$, $\forall j \in\left\{1 \ldots n_{b}\right\}$. Note that multiple actions could tie for the minimum.
Theorem 2. Regardless of how the first method breaks ties, the following holds $\forall b \in \mathbb{B}$ and $\forall t \in \mathbb{N}$. ${ }^{\ddagger}$

$$
\begin{equation*}
\sum_{i=1}^{n_{a}^{\star}} \min \left(X_{i, b}, 0\right) \leq \operatorname{Risk}_{b}(D, t) \leq \sum_{i=1}^{n_{a}^{\star}} \max \left(X_{i, b}, 0\right) \tag{8}
\end{equation*}
$$

Method 2 While the previous approach helps less-often chosen actions catch up to the others, the next approach chooses actions that - if used - will get the least ahead of the others. The sequence of decisions $D^{\prime}$ for this method is described formally as follows:

$$
\begin{equation*}
d^{\prime}(t) \in\left\{j \leq n_{a}^{\star} \mid \forall i \leq n_{a}^{\star}: \frac{k_{j}\left(D_{t-1}^{\prime}\right)+1}{F_{j}^{\star}} \leq \frac{k_{i}\left(D_{t-1}^{\prime}\right)+1}{F_{i}^{\star}}\right\} \tag{9}
\end{equation*}
$$

Theorem 3. Regardless of how the second method breaks ties, the following holds for $\forall b \in \mathbb{B}$ and $\forall t \in \mathbb{N}^{*}{ }^{\star}$

$$
\begin{equation*}
-\sum_{i=1}^{n_{a}^{\star}} \max \left(X_{i, b}, 0\right) \leq \operatorname{Risk}_{b}\left(D^{\prime}, t\right) \leq-\sum_{i=1}^{n_{a}^{\star}} \min \left(X_{i, b}, 0\right) . \tag{10}
\end{equation*}
$$

## Discussion

We proposed two methods for choosing actions such that the risks are always lower-bounded. While GR leximinoptimizes next-step risks, our methods greedily optimize actions' usage frequencies (relative counts) relative to some optimal configuration $F^{\star}$; we will call our methods GF.

We denote with $L B_{1}$ and $L B_{2}$ the lower bounds on risks guaranteed by the two GF methods respectively.

$$
\begin{align*}
& L B_{1}=\min _{b} \sum_{i=1}^{n_{a}^{\star}} \min \left(X_{i, b}, 0\right)  \tag{11}\\
& L B_{2}=\min _{b} \sum_{i=1}^{n_{a}^{\star}}-\max \left(X_{i, b}, 0\right) . \tag{12}
\end{align*}
$$

There is an obvious way to unify the two: compute $L B_{1}$ and $L B_{2}$ ahead of time, then use the most promising method. Thus our best worst risk guarantee is:

$$
\begin{equation*}
W R \geq \max \left(L B_{1}, L B_{2}\right) \tag{13}
\end{equation*}
$$

The computational complexity of our algorithms is $O\left(\lg n_{a}^{\star}\right)$ per time step. This is because one can use a heap to store actions' $k_{j}[+1] / F_{j}^{\star}$ scores, since our algorithms change a single action's score at each time step.

Eliminating Unnecessary Actions If the vector $F^{\star}$ is not unique, the particular choice of $F^{\star}$ influences both the time complexity (through $n_{a}^{\star}$ ), and the worst risk. Each beneficiary's worst risk is bounded below by the sum of his negative $X$ values (or the negative sum of his positive $X$ values), so eliminating an action can only improve the worst risk.
Lemma 1. There must always exist $F^{\star}$ such that $n_{a}^{\star} \leq n_{b}{ }^{\ddagger}$
Based on this result, one can eliminate at least $\max \left(n_{a}-\right.$ $\left.n_{b}, 0\right)$ actions in a preprocessing phase, thus reducing the per-step complexity to $O\left(\lg \left(\min \left(n_{a}, n_{b}\right)\right)\right)$. There exists a simple algorithm for this task based on a particular constructive proof for Carathéodory's theorem (e.g. (Florenzano and Le Van 2001)), using Gauss elimination. The complexity of the Gauss elimination is $O\left(n_{a} n_{b}^{2}\right)$, and there are at most $n_{a}-1$ iterations, so the entire pre-process of eliminating actions can be done in $O\left(n_{a}^{2} n_{b}^{2}\right)$.

Breaking Ties with GR The lower-bounds on worst risks (Equations 11 and 12) make no assumptions on how ties are broken, which means they assume the ties are broken in the worst possible way. Finding the best way to break ties is NPhard (see the proof of Theorem 1), so a greedy approach will have to do. The most obvious solution would be to use GR to break ties for GF (for a time complexity of $O\left(n_{b} \times n_{a}^{\star}\right)$ per time step).

## Related Work

The game-theoretic work in (Bhaskar 2000; Lau and Mui 2005) is concerned with finding Nash equilibria that result in alternating joint-actions (also referred to as turn-taking), but these results were tailored for specific classes of 2-by-2 games.

The starting point in our research on long-term fairness was the work in (Verbeeck, Parent, and Nowé 2003) on "periodic policies." Their reward model comes in the form of a normal-form game, but the players are actually cooperative learning agents (rather than self-interested). The process consists of the learners playing selfishly to discover a pure Nash equilibrium, while being interrupted periodically to compare accumulated rewards. The player gaining the most (in the current Nash equilibrium and overall) has its action put off-limits until the others catch up. Alternatively (Verbeeck et al. 2006), after the learners discover all pure Nash equilibria, they create a periodic policy consisting of those joint actions with the fairest outcomes. In the authors' examples the players have only two actions: a highly lucrative one and a social one they fall back on while waiting for the other(s) to catch up. Both of these greedy algorithms could lead to utility profiles arbitrarily worse than the optimum.

The Worst Risk Maximization Problem is actually a variant of the Compact Vector Summation Problem (Sevast' janov 1991): given a finite set of vectors $X_{1}, \ldots, X_{n} \in \mathbb{R}^{m}$ such that $\sum_{i=1}^{n} X_{i}=0$, one must find a permutation $\pi$ of $\{1,2, \ldots, n\}$ that minimizes $\max _{1 \leq k \leq n}\left\|\sum_{i=1}^{k} X_{\pi_{i}}\right\|$. In the earliest such work $\|\cdot\|$ was the Euclidian norm, so the problem consisted of ordering the vectors such that the path resulting from adding vectors one by one stays inside a minimum-radius $m$-dimensional circle centered at the origin. Later research has focused on results general enough to accommodate arbitrary norms (intuitively a norm is a function associating a "size" value to every vector).

We submit that the algorithm in (Sevast'janov 1991) is the most relevant algorithm for a comparison with our algorithms: it has the best guarantee and time complexity we are aware of and it accommodates "asymmetric norms." The last item is relevant because it is more meaningful to compare against algorithms trying to optimize the same function, and the function we try to optimize can be rearranged as an asymmetric norm but not a norm. Maximizing the worst risk is equivalent to minimizing is the largest absolute value of any negative coordinate of any partial sum of $X$ vectors. This function is an asymmetric norm, but not a norm, since it satisfies the triangle inequality $(\|y+z\| \leq\|y\|+\|z\|)$ and positive definiteness $(\|y\|=0 \Rightarrow y=0)$, but it only satisfies the scalability condition $(\|k y\|=|k| \times\|y\|)$ for positive scaling factors (Borodin 2001).

Guarantee Comparison Sevat'janov's algorithm guarantees no risk will be worse than $-M\left(n_{a}^{\star}-1+\frac{1}{n_{a}^{\star}}\right)$, where $M=\max _{1 \leq k \leq n}\left\|X_{i}\right\|$ (i.e. $-M$ is the smallest coordinate of any X vector size). We submit that $L B_{1} \geq-M\left(n_{a}^{\star}-1\right)$. This results from Equation 11, replacing every negative $X_{i, b}$ with $-M$, and noticing that for any given beneficiary $b$ at most
$n_{a}^{\star}-1$ of his $X$ values can be negative. ${ }^{\ddagger}$ Therefore our worst risk, $W R \geq L B_{1} \geq-M\left(n_{a}^{\star}-1\right)=-M\left(\min \left(n_{a}, n_{b}\right)-1\right)>$ $-M\left(n_{b}-1+\frac{1}{n_{b}}\right)$, the guarantee of Sevast'janov's algorithm.

Complexity Comparison Sevast'janov's algorithm has a complexity of $O\left(n^{2} m^{2}\right)$, because it picks a vector $n-m$ times, and each such operation has a complexity of $O\left(\mathrm{~km}^{2}\right)$, where $k$ is the number of alternatives $(k=n, \ldots, n-m)$. An iteration in Sevast' janov's algorithm has the same complexity as an iteration in our preprocessing phase (they are both based on Gauss elimination). We also submit that the number of iterations in Sevast'janov's algorithm is at least the number of iteration in our preprocessing phase (each action has at least multiplicity 1). Therefore the time complexity of our preprocessing phase cannot be larger than that of Sevastianov's algorithm. More importantly, even if that algorithm were extended to benefit from our preprocessing phase and to explicitly deal with multiplicities (i.e. $k=n_{a}^{\star}$ ), its complexity would still be $O\left(n_{a}^{\star} n_{b}^{2}\right)$ per time step which is higher than the complexity of our algorithms (even when breaking ties with GR).

In summary, by eliminating unnecessary actions and only keeping track of multiplicities, we are able to offer worst case guarantees that are never worse than (and sometimes much better than) Sevast'janov's.

## Conclusions

In this paper we studied the problem of achieving certain long-term fairness guarantees in a simple repeated-game setup: (1) all beneficiaries are entitled to their sociallyoptimal utilities and (2) no matter when the game ends all beneficiaries are guaranteed to have received close to what they were entitled to that point. We proved that finding an optimal solution with respect to the second guarantee is NPhard and proposed two efficient approximation algorithms.

## Acknowledgements

We thank Octav Olteanu, Joey Harrison, Zoran Duric, Alexei Samsonovich, and Alex Brodsky for their help.

## References

Balan, G.; Richards, D.; and Luke, S. 2008. Long-term fairness with bounded worst-case losses. Technical Report GMU-CS-TR-2008-2, Department of Computer Science, George Mason University.
Bhaskar, V. 2000. Egalitarianism and efficiency in repeated symmetric games. Games and Economic Behavior 32(2):247-262.
Borodin, P. A. 2001. The Banach-Mazur theorem for spaces with asymmetric norm. Mathematical Notes 69:298-305.
Florenzano, M., and Le Van, C. 2001. Finite Dimensional Convexity and Optimization. Springer.
Lau, S.-H. P., and Mui, V.-L. 2005. Using turn taking to mitigate conflict and coordination problems in the battle of the sexes game. Technical Report 1129, Hong Kong Institute of Economics and Business Strategy.

Potters, J. A. M., and Tijs, S. H. 1992. The nucleolus of a matrix game and other nucleoli. Mathematics of Operations Research 17(1):164-174.
Sevast'janov, S. 1991. On the compact summation of vectors. Diskretnaya Matematika 3(3):66-72.
Verbeeck, K.; Nowé, A.; Parent, J.; and Tuyls, K. 2006. Exploring selfish reinforcement learning in repeated games with stochastic rewards. Journal of Autonomous Agents and Multi-Agent Systems 14(3):239-269.
Verbeeck, K.; Parent, J.; and Nowé, A. 2003. Homo egualis reinforcement learning agents for load balancing. In Innovative Concepts for Agent-Based Systems: 1 ${ }^{\text {st }}$ International Workshop on Radical Agent Concepts, volume 2564 of Lecture Notes in Computer Science, 81-91.


[^0]:    Copyright (C) 2008, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

[^1]:    ${ }^{\ddagger}$ Due to limited space we omit all proofs and mathematical derivations. They may be found in (Balan, Richards, and Luke 2008), as well as an extended discussion.

