

Fair Division of Indivisible Goods under Risk

Charles Lumet¹ and Sylvain Bouveret² and Michel Lemaître³

Abstract. We consider the problem of fairly allocating a set of m indivisible objects to n agents having additive preferences over them. In this paper we propose an extension of this classical problem, where each object can possibly be in bad condition (e.g broken), in which case its actual value is zero. We assume that the central authority in charge of allocating the objects does not know beforehand the objects conditions, but only has probabilistic information. The aim of this work is to propose a formal model of this problem, to adapt some classical fairness criteria to this extended setting, and to introduce several approaches to compute optimal allocations for small instances as well as sub-optimal good allocations for real-world inspired allocation problems of realistic size.

1 INTRODUCTION

The problem of allocating a set of indivisible objects to a set of agents arises in a wide range of applications including auctions, divorce settlements, frequency allocation, airport traffic management, scheduling, courses allocation, multiagent resource allocation and resource-sharing in general. In such problems, one usually needs to find fair solutions, where fairness refers to the need for compromises between the agents often conflicting preferences for the objects.

In most fair division problems, it is assumed that objects conditions are perfectly known when agents value them. However, it often happens that the actual value or utility of an object for an agent is only discovered after the allocation has been decided. As a toy example, let the objects be bottles of wine to be shared amongst a set of agents. The wine in each bottle can be corked, but this is not known before opening (which happens only after the bottles have been allocated). How can this uncertainty, which also arises in numerous variants of the aforementioned real-world problems, be taken into account? In this paper we present a simple variant of the multiagent resource allocation problem (see e.g. [7] for a survey on this topic) in which the decision-maker has probabilistic information about the objects conditions, conditions which affect agent utilities.

Uncertainty (or, more precisely, risk) issues in collective decision making have been studied before, as in the microeconomical works by Myerson [18] and more recently Gajdos and Tallon [10]. However, these works focus on continuous allocation spaces (divisible “objects”) and leave aside computational issues. Risk management in (combinatorial) auctions has also been studied (for example in [12]), more precisely the influence of potential bid withdrawals in terms of loss of total revenue for the auctioneer; but to the best of our knowledge, fairness issues are not considered. Finally, uncertainty in vote (which is another kind of collective decision problem) has been investigated in some recent papers. Yet, in most of these, uncertainty

stands for the incomplete knowledge of the decider about the agents’ preferences (see e.g. [2]). One notable exception is the work by Lu and Boutilier [14] which deals with a voting problem where each candidate may withdraw from the election.

In this article, three main assumptions are made. (i) The allocation is *centralized*, that is, it is decided and computed by a central benevolent authority, according to the agents’ preferences. (ii) Each object can only be in two possible conditions: good or bad. The actual condition of an object is only known with a given probability when the allocation is decided, but is known for sure after the objects have been allocated. (iii) The agents have (subjective) additive utility preferences over the objects.

Even though this framework seems restrictive, we claim that it is worth studying it because first, additivity is very natural as soon as preferences over sets of objects have to be represented in a compact way (e.g, the well-known Santa-Claus fair division problem [3] is based on additive preferences); second, in many real-world problems, uncertainty can be defined “object-wise” and thus can be naturally modeled as we do; lastly, as we will see, our framework raises non trivial computational issues, despite its apparent simplicity.

This article is structured as follows. In Section 2, we describe our framework for fair division of indivisible objects under risk. In Section 3 we propose some extensions of classical fairness criteria based on collective utility functions and on fair share guarantee and explore their theoretical properties. Finally, we propose in Section 4 some practical methods to compute optimal and approximated fair allocations and we give some experimental results.

2 MODEL

In the following, we use lower case bold font to represent vectors and upper case bold font to represent matrices.

A finite set of indivisible *objects* $\mathcal{O} = \{1, \dots, m\}$ must be allocated by a central authority to a finite set of *agents* $\mathcal{A} = \{1, \dots, n\}$. An *allocation* is a vector of shares $\pi = \langle \pi_1, \dots, \pi_n \rangle$ where $\pi_i \subseteq \mathcal{O}$, and $j \in \pi_i$ if and only if object j has been given to agent i . The set of *feasible allocations* is $\mathcal{F} = \{\pi, i \neq i' \Rightarrow \pi_i \cap \pi_{i'} = \emptyset\}$ (an object cannot be given to more than one agent). We further denote by $\pi_0 = \mathcal{O} \setminus \bigcup_{i \in \mathcal{A}} \pi_i$ the set of non allocated objects, if any.

Each object can be either in *good* or *bad condition*. The objects’ conditions are known only after the allocation has been made, but the decision-maker is nevertheless given probabilistic information: a vector $\mathbf{p} \in [0; 1]^m$ maps each object j to its probability p_j to be in good condition; these probabilities are assumed to be independent. Let $\mathcal{S} = 2^m$ be the set of the possible states of nature, and $\text{good}(s)$ the set of objects in good condition when state of nature s happens. We thus have $\forall s \in \mathcal{S}, \Pr(s) = \prod_{j \in \text{good}(s)} p_j \prod_{j \notin \text{good}(s)} (1 - p_j)$.

The agents’ preferences for the objects are numerically expressed by weights: $w_{i,j} \in \mathbb{R}^+$ represents the utility of object j in good condition for agent i . A weight of zero means that the agent is

¹ Onera Toulouse, France, email: charles.lumet@onera.fr

² Grenoble INP – LJG, France, email: sylvain.bouveret@imag.fr

³ Formerly Onera Toulouse, email: lmtr@club-internet.fr

not interested in the object. For the agent i , the individual utility of the allocation π if the state of nature s happens is defined as $u_i(\pi, s) \stackrel{\text{def}}{=} \sum_{j \in \text{good}(s) \cap \pi_i} w_{i,j}$. In other words, we assume that agents have *additive* preferences over the objects: an agent individual utility for an allocation and a state of nature is given by the sum of the weights of the objects in good condition received by said agent and each object in bad condition brings no extra utility.

To sum up, an instance of our problem is given by a tuple $\mathcal{P} = (\mathcal{A}, \mathcal{O}, \mathbf{p}, \mathbf{W})$, where $\mathcal{A} = \{1, \dots, n\}$ is a set of agents, $\mathcal{O} = \{1, \dots, m\}$ is a set of objects, $\mathbf{p} \in [0; 1]^m$ expresses the probability for each object to be in good condition, and $\mathbf{W}$ is the n -lines m -columns matrix of weights given to the objects by the agents.

Example 1 Table 1 shows an instance of our problem, with the probabilities of each possible state of nature (line (2)) and utility profiles associated with a given allocation (lines (3) and (4)).

3 FAIRNESS CRITERIA UNDER RISK

3.1 Ex-post vs ex-ante collective utilities

A classical way to define the quality of an allocation is to *aggregate* the agents' individual utilities with a commutative and increasing *collective utility function* (CUF) $G : (\mathbb{R}^+)^n \rightarrow \mathbb{R}^+$, which measures social welfare. A general survey on CUF can be found in [17]. In this paper, we will focus on the following ones:

- utilitarian and egalitarian CUF: $G^u \stackrel{\text{def}}{=} \sum$ and $G^e \stackrel{\text{def}}{=} \min$;
- Nash CUF: $G^n \stackrel{\text{def}}{=} \prod$;
- Ordered Weighted Averages (OWAs) [21]: $G^\omega(\mathbf{u}) \stackrel{\text{def}}{=} \sum_{i=1}^n \omega_i u_i^\uparrow$, where $\sum_{i=1}^n \omega_i = 1$, $\omega_i \geq 0 \forall i$, and $\mathbf{u}^\uparrow$ is a permutation of $\mathbf{u}$ s.t. $u_1^\uparrow \leq \dots \leq u_n^\uparrow$;
- Sums of Powers⁴ (SPs) [17]: $G^p(\mathbf{u}) \stackrel{\text{def}}{=} \text{sign}(p) \sum_{i=1}^n u_i^p$ for $p \neq 0$, and $G^0(\mathbf{u}) \stackrel{\text{def}}{=} \sum_{i=1}^n \log(u_i)$.

OWAs and SPs provide two families of compromises between the utilitarian and egalitarian CUF: the first one can be represented⁵ by $G^{(1/n, 1/n, \dots, 1/n)}$ or G^1 and the second one by $G^{(1, 0, \dots, 0)}$ or G^p , with $p \rightarrow -\infty$ (more precisely, this last aggregation function represents the leximin preorder, which refines the egalitarian CUF). Moreover, the Nash CUF can be represented by G^0 . We must also note that most of these operators only make sense if the individual utilities are *normalized* (in our case, imposing $\sum_{j \in \mathcal{O}} w_{i,j} = \sum_{j \in \mathcal{O}} w_{i',j} \forall i, i'$ is a reasonable way to do it – see e.g. [17] for other normalization methods). In the following we will write $G_{i \in \mathcal{A}} u_i$ for $G(\mathbf{u})$.

Since the actual state of nature is unknown before the allocation is chosen, utilities must be aggregated in all different possible states of nature. For this purpose we use the expected utility – other choices are possible, such as the min (pessimistic) or max (optimistic) operators, but we stand here on the most natural way of doing this aggregation. Depending on whether aggregation is first made over agents and then over states of nature or the other way around, we obtain two different functions [11, 18]: $acu : \mathcal{F} \rightarrow \mathbb{R}^+$, defined in (1), is called *ex-ante* collective utility and $pcu : \mathcal{F} \rightarrow \mathbb{R}^+$, defined in (2) is called *ex-post* collective utility.

$$\forall \pi \in \mathcal{F}, \quad acu(\pi) \stackrel{\text{def}}{=} G_{i \in \mathcal{A}} \left(\sum_{s \in \mathcal{S}} \Pr(s) \cdot u_i(\pi, s) \right) \quad (1)$$

$$\forall \pi \in \mathcal{F}, \quad pcu(\pi) \stackrel{\text{def}}{=} \sum_{s \in \mathcal{S}} \Pr(s) \cdot (G_{i \in \mathcal{A}} u_i(\pi, s)) \quad (2)$$

⁴ Note that this family corresponds (up to a p^{th} -root) to the family of root-power quasi-arithmetic means studied e.g. in multicriteria decision aid [15].

⁵ f represents g means formally that $\forall (x, y), f(x) \leq f(y) \Leftrightarrow g(x) \leq g(y)$.

Defining some *expected weights* $\tilde{w}_{i,j} \stackrel{\text{def}}{=} p_j w_{i,j}$, Proposition 1 shows that the *ex-ante* collective utility can be written as the collective utility of a risk-free equivalent instance (proof is easy thus omitted).

Proposition 1 Let $\tilde{u}_i(\pi) \stackrel{\text{def}}{=} \sum_{j \in \pi_i} \tilde{w}_{i,j}$ be the expected utility of agent i for the allocation π . Then $\forall \pi \in \mathcal{F}$, $acu(\pi) = G_{i \in \mathcal{A}} \tilde{u}_i(\pi)$.

3.2 Properties of ex-post and ex-ante utilities

We might wonder if we can find any relation between *ex-ante* and *ex-post* collective utilities. Actually, the following (in)equalities follow immediately from the definition of convexity, linearity, concavity:

Proposition 2 Let G be a CUF. Then $\forall \pi$, $pcu(\pi) \leq acu(\pi)$ (resp. $\geq, =$) if and only if G is concave (resp. convex, linear).

This proposition shows that determining whether the *ex-post* utility is greater than the *ex-ante* utility comes down to analyzing the concavity of the collective utility function, which is not completely straightforward for standard operators. We will give here several results about the classical CUF that favor equity in the solution, more formally, on those that promote *Pigou-Dalton transfers* [8, 19], that is, utility transfers from wealthier agents to poorer ones.

Definition 1 (Pigou-Dalton transfer) Let $\mathbf{u}$ and $\mathbf{u}'$ be two utility vectors. $\mathbf{u}'$ is obtained by a Pigou-Dalton transfer from $\mathbf{u}$ if and only if one can find a couple of distinct agents (i_1, i_2) , $i_1 \neq i_2$ such that (i) the sum of their utilities remains unchanged: $u_{i_1} + u_{i_2} = u'_{i_1} + u'_{i_2}$ (ii) inequalities decrease: $|u'_{i_1} - u'_{i_2}| < |u_{i_1} - u_{i_2}|$ (iii) other agents utilities remain unchanged: $\forall i \in \mathcal{A} \setminus \{i_1, i_2\}, u_i = u'_i$.

A CUF is said to *reduce inequalities* if for all $\mathbf{u}, \mathbf{u}'$, $G(\mathbf{u}') \geq G(\mathbf{u})$ as soon as $\mathbf{u}'$ is obtained by Pigou-Dalton transfer from $\mathbf{u}$. Note that this property is also called *Schur-concavity* in the theory of majorization (see e.g. [16]).

We will now show that reducing inequalities entails concavity for a general set of CUF, namely the *separable* ones (see e.g. [17]).

Definition 2 (Separability) A CUF G is said to be separable if and only if it exists an increasing function $g : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ s.t., for each utility vector $\mathbf{u} \in (\mathbb{R}^+)^n$, $G(\mathbf{u}) = \sum_{i=1}^n g(u_i)$.

Proposition 3 Let G be a separable CUF reducing inequalities. Then, G is concave, that is, the following inequality stands:

$$\forall \pi \in \mathcal{F}, \quad pcu(\pi) \leq acu(\pi) \quad (3)$$

The proof of this proposition follows rather directly from Jensen's inequality. Now, even if a wide range of CUF are separable (for example the sums of powers), it is not the case for the OWA. However, it appears that any OWA reducing inequalities is concave :

Proposition 4 Let G^ω be an OWA reducing inequalities. Then G^ω is concave, that is, Inequality (3) also stands.

We prove this using two lemmas (the proof of Lemma 1 is easy, thus omitted due to lack of space).

Lemma 1 G^ω reduces inequalities if and only if $\omega_1 \geq \dots \geq \omega_n$.

Lemma 2 Let $\mathbf{u} = (u_1, \dots, u_n)$ a vector, σ a permutation of $\{1, \dots, n\}$ s.t. $\forall i < j, u_{\sigma(i)} \leq u_{\sigma(j)}$, and ω a vector of sorted weights: $\omega_1 \geq \omega_2 \geq \dots \geq \omega_n$. Then $\sum_{i=1}^n (\omega_{\sigma(i)} - \omega_i) u_{\sigma(i)} \geq 0$.

Table 1. Utility profile and *ex-ante* and *ex-post* utility computation for a problem with 2 agents, 4 objects, probabilities $\mathbf{p} = \langle 0.8, 0.8, 0.5, 0.2 \rangle$, weights $\mathbf{w}_1 = \langle 10, 2, 4, 7 \rangle$ and $\mathbf{w}_2 = \langle 3, 8, 4, 10 \rangle$, allocation $\pi = \langle \{1, 4\}, \{2, 3\} \rangle$, $G = \min$. Here, $pcu(\pi) = 4.66$ and $acu(\pi) = 8.4$ (see Section 3.1).

(1)	$good(s)$	$\emptyset$	$\{1\}$	$\{2\}$	$\{3\}$	$\{4\}$	$\{1, 2\}$	$\{1, 3\}$	...	$\{2, 3, 4\}$	$\{1, 2, 3, 4\}$	expected utility
(2)	$Pr(s)$	0.016	0.064	0.064	0.016	0.004	0.256	0.064	...	0.016	0.064	
(3)	$u_1(\pi, s)$	0	10	0	0	7	10	10	...	7	17	9.4
(4)	$u_2(\pi, s)$	0	0	8	4	0	8	4	...	12	12	8.4
(5)	collective utility	0	0	0	0	0	8	4	...	7	12	4.66

Proof We can prove the result by induction on n . The base case ($n = 1$) is immediate: σ is the identity and the sum equals 0.

Suppose that the lemma is true for every vector $\mathbf{u}$ and ω of dimension $n - 1$. Let $\mathbf{u}$ and ω be two vectors of dimension n satisfying the conditions of Lemma 2, and let σ be a permutation sorting the components of $\mathbf{u}$. Let $t = \sigma^{-1}(n)$ (thus $n = \sigma(t)$). The following inequalities stand: $\omega_t \geq \omega_n(4)$ and $u_{\sigma(t)} = u_n \leq u_{\sigma(n)}(5)$.

Focusing on the sum components that contain t and n , we obtain:

$$\begin{aligned} & (\omega_{\sigma(t)} - \omega_t)u_{\sigma(t)} + (\omega_{\sigma(n)} - \omega_n)u_{\sigma(n)} \\ &= (\omega_{\sigma(n)} - \omega_t)u_{\sigma(n)} + \underbrace{[(\omega_t - \omega_n)(u_{\sigma(n)} - u_n)]}_{\geq 0 \text{ (from inequalities (4) and (5))}} \end{aligned} \quad (6)$$

Now let σ' be the permutation of $\{1, \dots, n - 1\}$ defined by:

$$\sigma' : i \mapsto \begin{cases} \sigma(n) & \text{if } i = t \\ \sigma(i) & \text{otherwise} \end{cases}$$

We can prove, by using Equation (6), definition of σ , and induction hypothesis:

$$\sum_{i=1}^n (\omega_{\sigma(i)} - \omega_i)u_{\sigma(i)} \geq \sum_{i=1}^{n-1} (\omega_{\sigma'(i)} - \omega_i)u_{\sigma'(i)} \geq 0,$$

which proves the hypothesis at rank n , and completes the proof. ■

Proof (Proposition 4) Let G^ω be an OWA that reduces inequalities. By Lemma 1, $\omega_1 \geq \omega_2 \geq \dots \geq \omega_n$. Let π be an allocation, and as in Proposition 1, $\tilde{u}_i(\pi) = \sum_{s \in \mathcal{S}} Pr(s)u_i(\pi, s)$ agent i 's expected utility. Suppose w.l.o.g that $\tilde{u}_1(\pi) \leq \tilde{u}_2(\pi) \leq \dots \leq \tilde{u}_n(\pi)$ (we can permute the agents if it is not the case). For all $s \in \mathcal{S}$ let σ_s be a permutation such that $\forall i < i', u_{\sigma_s(i)}(\pi, s) \leq u_{\sigma_s(i')}(\pi, s)$. We have $acu(\pi) = \sum_{i=1}^n \sum_{s \in \mathcal{S}} Pr(s)\omega_i u_i(\pi, s)$, and $pcu(\pi) = \sum_{i=1}^n \sum_{s \in \mathcal{S}} Pr(s)\omega_i u_{\sigma_s(i)}(\pi, s)$. Thus:

$$acu(\pi) - pcu(\pi) = \sum_{s \in \mathcal{S}} Pr(s) \times \sum_{i=1}^n ((\omega_{\sigma_s(i)} - \omega_i) \cdot u_{\sigma_s(i)}(\pi, s))$$

Applying Lemma 2 on vectors $\mathbf{u}(\pi, s)$ and ω , and on permutation σ_s , one obtains immediately $acu(\pi) - pcu(\pi) \geq 0$. ■

We can notice that Propositions 3 and 4 immediately entail that the inequality $pcu \leq acu$ holds for the egalitarian and utilitarian CUF (actually $pcu = acu$ in the utilitarian case).

The Nash CUF case kind of lies in-between: if we stick to the original definition ($G = \prod$), the inequality $pcu \leq acu$ does not hold. Consider for example the two vectors $\mathbf{x} = (1, 1)$ and $\mathbf{y} = (3, 3)$. We have $0.5 \times (1 \times 1 + 3 \times 3) = 5 \geq (0.5 \times (1 + 3)) \times (0.5 \times (1 + 3)) = 4$, which proves this CUF is not concave (it can be shown that it is not convex either). Yet, if we consider the representation of the Nash CUF as a SP operator G^0 , we have proved earlier its concavity, and thus the inequality $pcu \leq acu$ holds. This example shows that the latter property is not inherent to a social welfare ordering, but rather to its chosen numerical representation.

3.3 Fair share guarantee under risk

The optimization of CUF focuses on the search for the best allocations (for a particular meaning of what “best” is), which can be very demanding. Another approach is to look for allocations that satisfy a given decision criterion. One prominent approach, defended by [20], is based on the *fair share principle* which is defined as follows in a risk-free setting:

Definition 3 An allocation π satisfies the fair share test if and only if for all i , $u_i(\pi) \geq 1/n \sum_{j=1}^m w_{i,j}$.

The fair share test is also known as proportionality principle, since it is based on the fact that each agent considers that she is entitled to at least $1/n$ of what she considers to be the total value of the resource to share. One appealing property of this test is its independence on the individual utility scales, so there is no need to bother about normalization of individual utilities.

The fair share test can be extended in several ways to risk. The *ex-ante* fair share test is a straightforward extension:

Definition 4 An allocation π satisfies the *ex-ante* fair share test if and only if for all i , $\tilde{u}_i(\pi) \geq 1/n \sum_{j=1}^m \tilde{w}_{i,j}$.

In other words, ensuring *ex-ante* fair share means that each agent is considered happy as soon as she gets at least her entitlement of *expected* utility. Now, applying the same argument to define *ex-post* fair share leads to the following criterion: an allocation passes the *ex-post* fair share test if and only if it passes the fair share test in *every* state of the world. As soon as for a given state (may it have the tiniest non zero probability) an agent only gets objects in bad state, while the other ones are in good state, the allocation will not pass the test. Hence, for most instances, no allocation will pass the test, which argues in favor of a less demanding criterion. A possibility is to maximize the probability to satisfy fair share, instead of trying to ensure it in each state of the world.

Definition 5 The *ex-post* probability of fair share of a given allocation π is defined as follows:

$$\varphi^p(\pi) \stackrel{\text{def}}{=} \sum_{s \in \mathcal{S}} Pr(s) \min_{i \in \mathcal{A}} (\varphi_i(\pi, s)),$$

with $\varphi_i(\pi, s) = 1$ if $u_i(\pi, s) \geq 1/n \sum_{j \in good(s)} w_{i,j}$, 0 otherwise.

The value $\varphi_i(\pi, s)$ indicates whether agent i has her fair share in state s . The use of operator $\min$ comes from the definition of fair share test which is satisfied only if all the agents have their fair share.

We can notice that since $\varphi_i(\pi, s)$ is a numerical value, depending on a pair (agent, state), we can use it exactly the same way we used individual utilities. Therefore, since we have defined an *ex-post* probability of fair share, we can also define the *ex-ante* version of this criterion⁶:

⁶ Also note that we could replace the $\min$ operator with another one (for example $\sum$), which would lead to other criteria, like *ex-ante* or *ex-post* quantities of fair share for example.

Definition 6 The *ex-ante* probability of fair share of a given allocation π is defined as follows:

$$\varphi^a(\pi) \stackrel{\text{def}}{=} \min_{i \in \mathcal{A}} \sum_{s \in \mathcal{S}} \Pr(s) \varphi_i(\pi, s).$$

We might wonder if there are some links between the definition of *ex-ante* probability of fair share (Definition 6) and the *ex-ante* fair share test (Definition 4). It turns out that even if $(\varphi^a = 1)$ entails passing the *ex-ante* fair share test, passing the *ex-ante* fair share test does not guarantee a higher value of φ^a :

Proposition 5 For all allocation π , $(\varphi^a(\pi) = 1) \Rightarrow \pi$ satisfies the *ex-ante* fair share test. However, for any pair (π, π') , $(\pi'$ satisfies the *ex-ante* fair share test and π does not) $\nRightarrow \varphi^a(\pi') > \varphi^a(\pi)$.

Proof Let π be such that $\varphi^a(\pi) = 1$. By Definition 6, $\forall(i, s)$, $\varphi_i(\pi, s) = 1$, which entails $u_i(\pi, s) \geq 1/n \sum_{j \in \text{good}(s)} w_{i,j}$. Then, $\tilde{u}_i(\pi) = \sum_{s \in \mathcal{S}} \Pr(s) u_i(\pi, s) \geq \sum_{s \in \mathcal{S}} \Pr(s) \times 1/n \sum_{j \in \text{good}(s)} w_{i,j} = 1/n \sum_{j=1}^m \tilde{w}_{i,j}$, which proves the first implication.

Now consider the following example:

$$\mathbf{W} = \begin{array}{c|cc} & j_1 & j_2 \\ \hline i_1 & 899 & 101 \\ \hline i_2 & 991 & 9 \end{array} \quad \begin{array}{l} \mathbf{p} = \langle 0.1, 0.9 \rangle \\ \pi = \langle \{j_2\}, \{j_1\} \rangle \\ \pi' = \langle \{j_1\}, \{j_2\} \rangle \end{array}$$

We have $\varphi^a(\pi') = \varphi^a(\pi) = 0.19$. π passes the *ex-ante* fair-share test, but π' does not. ■

The following result comes directly from the concavity of min:

Proposition 6 For each $\pi \in \mathcal{F}$, we have $\varphi^p(\pi) \leq \varphi^a(\pi)$.

Note also that both φ^a and φ^p are bounded above by 1, $\varphi^a = 1 \Leftrightarrow \varphi^p = 1$. It corresponds to the case where the fair share test is verified in each state of the world.

Finally, it can be noticed that in the risk-free case there is an interesting correlation between egalitarianism and fair share: if the set of allocations passing the fair share test is non empty, every optimal allocation will pass it. We will show that it is still the case for *ex-ante* fair share test and egalitarianism, but not for φ^p and egalitarianism: maximizing *ex-post* egalitarian CUF may lead to sub-optimal allocations concerning *ex-post* probability of fair share.

Proposition 7 Let $(\mathcal{A}, \mathcal{O}, \mathbf{p}, \mathbf{W})$ be an instance of the allocation problem, where the weight matrix is normalized in expectation (i.e. $\forall i, i', \sum_{j \in \mathcal{O}} \tilde{w}_{i,j} = \sum_{j \in \mathcal{O}} \tilde{w}_{i',j}$). If the set of allocations passing the *ex-ante* fair share test is non empty, then all the *ex-ante* egalitarian optimal allocations will pass it. However, the intersection between the set of *ex-post* egalitarian optimal allocations and the set of allocations maximizing φ^p may be empty.

Proof The proof concerning the *ex-ante* case is a direct extension of the result in the risk-free setting. Concerning the *ex-post* case, we can consider the following instance and allocations :

$$\mathbf{W} = \begin{array}{c|ccc} & j_1 & j_2 & j_3 \\ \hline i_1 & 6 & 2 & 2 \\ \hline i_2 & 4 & 1 & 5 \end{array} \quad \begin{array}{l} \mathbf{p} = \langle 0.9, 0.5, 0.4 \rangle \\ \pi = \langle \{j_1, j_2\}, \{j_3\} \rangle \\ \pi' = \langle \{j_1\}, \{j_2, j_3\} \rangle \end{array}$$

It can be proved that π' is the only one *ex-post* egalitarian optimal allocation. We have $\text{pcu}(\pi) = 1.84$, $\varphi^p(\pi) = 0.41$, $\text{pcu}(\pi') = 2.25$ and $\varphi^p(\pi') = 0.39$, which gives the counterexample needed. ■

4 COMPUTING FAIR ALLOCATIONS

In this section, we will introduce several approaches to compute fair allocations. Here, “fairness” will be interpreted in the sense of egalitarianism. Moreover, we will focus on *ex-post* egalitarianism, because (i) the choice of *ex-post* egalitarianism has been advocated by several authors, among which Adler and Sanchirico [1], Broome [5], and Fleurbaey [9] and (ii) there is a very natural interpretation of the *ex-post* approach, as the decision a Bayesian decision-maker would naturally take if she were paid *after* uncertainty has been resolved, according to the (certain) collective utility of her decision.

4.1 Experimental setting

To test our algorithms, we will focus on a particular class of real-world problems, time-sharing problems, in which the use of a common resource, owned by several agents, has to be divided in time. Managing a constellation of Earth observing satellites owned by several countries [13], or planning the use of a telescope shared between world astronomers, are two examples of this class of problems.

In such problems, the agents’ interests in using the resource might vary in time: for example, a country representative might be more interested to use an Earth observing satellite when it is located right above her country than when it is above the ocean. Moreover, the actual utility received by an agent for using the facility at a given time slot might depend on exogenous conditions which are not known beforehand : it could be for example the weather conditions which can seriously affect the quality of the satellite’s observations.

Assuming that time is divided into a finite number of indivisible and non sharable time slots, the framework introduced in Section 2 is general enough to encompass this class of problems. Real-world instances of this kind share two typical features: (i) the agents’ preferences are somewhat but not completely similar, that is, some objects will be fervently desired by all the agents, whereas some can be of interest for only one agent. (ii) contiguous time slots are strongly linked, which entails that they will likely share similar interest or probability of good condition⁷.

Algorithms introduced in this article are implemented using Java and run on random instances reproducing the features of aforementioned real-world problems (with normalized weights). Note that real-world instances typically involve a small number of agents (between 2 and 5) and a large number of objects (about 100).

4.2 Exact algorithms

Computing *ex-post* utility Before trying to compute an *optimal* allocation, we must be able to compute the collective utility for any given allocation. A naive algorithm for *computing* the *ex-post* collective utility, directly applying (2), would require the computation of the collective utility in each possible state (i.e. each column in Table 1), that is, the enumeration of an exponential number of values. For this reason we strongly suspect that computing the *ex-post* egalitarian collective utility of a given allocation is $\#P$ -complete (which would mean that computing an optimal allocation is even harder). However, the precise complexity of this problem remains open.

The algorithm we use improves a naive branching algorithm for computing the *ex-post* collective utility of an allocation by using the

⁷ Note that for the sake of simplicity, we stick to the probabilistic independence hypothesis, but we could also think of relaxing it and assume that probability variables of contiguous time slots are in fact dependent.

fact that as soon as all the objects allocated to an agent are in bad condition, the utility of this agent is zero, and so is the collective utility, whatever conditions the remaining objects are in. In order to quickly “eliminate” such states of nature, whose enumeration is unnecessary, the algorithm browses the states of nature by fixing first the state of objects allocated to agents whose share contains few objects.

Computing optimal ex-post allocations The *optimization* problem is tackled with a branch-and-bound algorithm. A basic upper bound $\overline{acu}(\pi)$ can be defined for each (partial) allocation π as the *ex-ante* utility of a (unfeasible) allocation which would allocate to all the agents the set π_0 of objects that are currently not allocated:

$$\overline{acu}(\pi) \stackrel{\text{def}}{=} \min_{i \in \mathcal{A}} (\sum_{j \in \pi_i} \tilde{w}_{i,j} + \sum_{j \in \pi_0} \tilde{w}_{i,j}).$$

$\overline{acu}$ only gives a rough approximation of pcu . We can provide a better upper bound by using a function which lies between pcu and acu . The idea is to compute utility in an *ex-ante* manner for a subset $\mathcal{O}_{ea}$ of objects, and in an *ex-post* manner for the other ones (those in $\mathcal{O}_{ep} = \mathcal{O} \setminus \mathcal{O}_{ea}$). In other words, we use *mixed individual utilities*:

$$mu_i(\pi, s, \mathcal{O}_{ea}) \stackrel{\text{def}}{=} \sum_{\substack{j \in \mathcal{O}_{ep} \cap \text{good}(s) \\ j \in \pi_i}} w_{i,j} + \sum_{\substack{j \in \mathcal{O}_{ea} \\ j \in \pi_i}} \tilde{w}_{i,j} \quad (7)$$

The *mixed collective utility* is then defined as the *ex-post* collective utility from mixed individual utilities, i.e by replacing $u_i(\pi, s)$ by $mu_i(\pi, s, \mathcal{O}_{ea})$ in Equation 2. We have the following proposition:

Proposition 8 For each concave CUF G , $\forall \pi \in \mathcal{F}, \forall \mathcal{O}_{ea} \subseteq \mathcal{O}$, $pcu(\pi) \leq mcu(\pi, \mathcal{O}_{ea}) \leq acu(\pi)$. Moreover, if $G = \min$ then $0 \leq mcu(\pi, \mathcal{O}_{ea}) - pcu(\pi) \leq \sum_{j \in \mathcal{O}_{ea}} \bar{p}_j p_j w_{a(j)j}$, where $\bar{p}_j \stackrel{\text{def}}{=} 1 - p_j$, and $a(j)$ is the agent receiving j .

The proof is omitted due to lack of space. This proposition implies that (i) we can use $mcu(\pi, \mathcal{O}_{ea})$ as an upper bound of $pcu(\pi)$, and (ii) if $G = \min$ this bound is tighter if $\mathcal{O}_{ea}$ is built for an allocation π with objects minimizing $p_j \bar{p}_j w_{a(j)j}$.

We can also notice that the mixed utility only depends on the condition of objects in $\mathcal{O}_{ep}$. Hence, the number of states in the main summation of mcu is halved for each object added inside $\mathcal{O}_{ea}$, property of which algorithms will take profit.

As said earlier, our optimization algorithm is based on a branch-and-bound approach, where nodes in the search space are the successive objects to allocate, and branches are the agents receiving them. Our algorithm is tailored for our fair division problem as follows.

- (i) We use the upper bound function $\overline{acu}$ to prune branches of the search tree during its exploration.
- (ii) We use dynamic heuristics: when selecting a new object to be allocated, the one preferred by the currently poorest (in expected utility) agent is chosen among those still left; each object will be firstly allocated to this poorest agent.
- (iii) Since the computation of mcu is too time-consuming to be used as an efficient pruning mechanism, we use this bound with profit as a last possible cut when a complete allocation has been found, to avoid an unnecessary (and costly) *ex-post* collective utility computation. The size of $\mathcal{O}_{ea}$ has been empirically fixed to $m/3$, which seems to be a good compromise between precision and computation time for the instances we tackle, and the objects to put in $\mathcal{O}_{ea}$ are chosen according to Proposition 8.

Results Table 2 shows the results for two versions of the branch-and-bound algorithm: (a) only uses $\overline{acu}$ as an upper bound (i), where (b) also involves dynamic heuristics (ii) and mcu (iii). The results show the efficiency of the latter, in which, among other, mixed collective utility dramatically reduces the number of *ex-post* collective utility computations.

Table 2. Exact resolution. Number of instances solved in 30 seconds (over 100 random instances) for different number of agents and objects.

n	m	(a)	(b)	n	m	(a)	(b)
2	≤ 15	100	100	3	≤ 11	100	100
2	20	0	100	3	16	0	97
2	25	0	60	3	21	0	48
2	30	0	1	3	26	0	7

4.3 Sub-optimal algorithms

As we can see *e.g* in Table 2, even with the improvements introduced, we remain unable to solve a majority of instances with 3 agents and 26 objects within 30 seconds with our exact algorithm (it typically requires between 5 and 10 minutes for this instance size). Moreover, even if we did not prove it formally, it is very likely that the mere computation of the *ex-post* egalitarian utility of a given allocation will be out-of-reach for more than a few dozens of objects, because of the exponential number of states of nature. The only reasonable way of solving instances of realistic size thus seems to use sub-optimal algorithms, for both the computation of a good allocation and the evaluation of the *ex-post* utility.

We can notice however that if we stick to the egalitarian *ex-post* criterion, we will be unable to evaluate the quality of the solutions, since exact computation of an optimal solution is out of reach. To overcome this issue and give a rough idea of how good our solutions are, we chose to switch our criterion to the probability of fair share φ^p (Definition 5), because this criterion has an obvious upper bound ($\varphi^p \leq 1$), and thus conveys in itself an information about the quality of the solution.

Estimating ex-post criterion The approximate computations are made according to the Monte-Carlo method: since browsing all the states of nature to get an exact result is impossible, we select a given number q of random states (according to their actual probability of occurrence) and check in each of them if the fair share is guaranteed.

Let X be a random variable, which takes its values in $\mathcal{S}$ according to the probability $\Pr(X = s) = \Pr(s)$. For a given $\pi \in \mathcal{F}$, let $\hat{\varphi}(X) = \min_{i \in \mathcal{A}} \varphi_i(\pi, X)$, which is an unbiased estimator of $\varphi^p(\pi)$. The Monte-Carlo method produces q independent draws, i.e. a sample of values $\hat{\varphi}_1, \hat{\varphi}_2, \dots, \hat{\varphi}_q$. A better estimator (with a smaller variance) of $\varphi^p(\pi)$ is then defined by $\hat{\varphi}' = 1/q \sum_{k=1}^q \hat{\varphi}_k$. Now, given a probability α , we can use the central limit theorem to give a confidence interval around which we have a probability $1 - \alpha$ to find the actual value of φ^p :

$$\Pr[|\hat{\varphi}' - \varphi^p| \geq C(\alpha)] \leq \alpha, \text{ with } C(\alpha) = F(\alpha) \sqrt{\frac{\text{Var}(\hat{\varphi})}{q}} \quad (8)$$

where F is the cumulative density function of a normal distribution of expectation $E(\hat{\varphi})$ and of standard deviation $\sigma(\hat{\varphi})$, and where $\text{Var}(\hat{\varphi})$ is the variance of $\hat{\varphi}$. Computing $\text{Var}(\hat{\varphi})$ requires browsing all the states of nature: thus, we use the samples we have to estimate it, which is common practice in statistics theory. Then, Equation (8) directly gives a confidence interval for our Monte-Carlo estimation.

Computing fair allocations Our algorithm follows the “Heuristic-Biased Stochastic Sampling” paradigm [4]. Complete allocations are built object by object, using the same heuristics as for the exact algorithms (see Section 4.2), but with a random bias, so that the search space is partially explored in the vicinity of the pure heuristics. When a complete allocation π has been built, an estimation of $\varphi^p(\pi)$ is made, with q_1 draws, in order to decide whether the allocation will




Figure 1. Evolution in time of the value of $\hat{\varphi}'$ during execution of the stochastic greedy algorithm, averaged over 100 instances, for $n = 3$, $m = 100$, with $q_1 = 200$, $q_2 = 5 \times 10^5$, $nb_S = 10$ and $nb_C = 50$.

be stored or not. A fixed number nb_S of such promising allocations is stored within the course of the algorithm; if an allocation is better – as far as this first approximate computation can tell – than the worst currently stored, it is saved. As soon as a total of nb_C allocations have been generated this way, a more precise estimation is made for each of the nb_S stored allocations using q_2 draws ($q_2 \gg q_1$), and only the best one is kept. The building of allocations then continues, until a given time has elapsed.

Parameters values have been chosen after some preliminary experiments, and give good results for the typical case $n = 3$, $m = 100$.

Results The approximate algorithm is tested on 100 instances, for a duration of 120 seconds each. Figure 1 illustrates the evolution of the best solution found, throughout the execution of the algorithm. Table 3 shows the data about the precision of this method.

m	10	20	50	100
exact variance of $\hat{\varphi}$	.1532	.1892	-	-
estimated variance of $\hat{\varphi}$, $q = 200$	.1702	.1904	.1141	.0200
estimated variance of $\hat{\varphi}$, $q = 10^5$	.1692	.1916	.0353	.0115
$C(1\%)$, $q = 200$	.0445	.0493	.1146	.0193
$C(1\%)$, $q = 5 \times 10^5$	.0014	.0016	.0011	.0004

Table 3. Estimation quality for $n = 3$, averaged over 100 instances.

As we can see in Table 3, on small instances the estimated variance is very close from the exact one. Regarding instances of realistic size, not only Figure 1 shows that our stochastic greedy approach quickly finds allocations whose estimated probability of fair share is very close to 1, but also Table 3 clearly indicates that this estimation is very likely to be close to the exact value. What experiments also show (but which does not appear in Figure 1 or Table 3) is that combining storage with two different numbers of draws (q_1 and q_2) allows for a quicker convergence than a simple greedy algorithm with q_2 draws, with a more precise confidence interval than an algorithm with q_1 draws.

In summary, these results clearly show that combining Monte-Carlo draws with our stochastic greedy algorithms can be used to solve efficiently real-world instances by providing (i) very good allocations in a short amount of time and (ii) precise estimates of φ^P .

5 CONCLUSION

In this paper we have introduced a formal model of the problem of fairly allocating a set of indivisible objects (goods, time-slots ...) when the conditions (and hence the utility for agents) of these objects are subject to uncertainty, as well as algorithms to solve it.

As pointed for a long time by economists, in such a problem involving uncertainty, two points of view can be taken in order to socially measure the possible allocations: either consider the social welfare derived from expected individual utilities (*ex-ante* approach), or consider the expected social welfare (*ex-post* approach). When fairness is a concern, many arguments are in favour of the last one.

We have investigated several properties of *ex-ante* and *ex-post* criteria, and introduced exact and sub-optimal algorithms to find allocations that maximize the *ex-post* egalitarian collective utility, or maximize the *ex-post* probability of fair share. As we have seen, the complexity makes the exact computation unfeasible in practice. However, we have managed to solve real-size instances with a remarkable level of precision by combining a stochastic greedy algorithm with Monte-Carlo estimations.

Worthwhile possible extensions of this work include the study of other fairness criteria, such as the maximin fair share, introduced in [6], which suits well to indivisible goods but requires solving numerous fair share subproblems. We also plan on designing other approximate evaluation methods to use in the stochastic greedy algorithm.

REFERENCES

- [1] M.D. Adler and C.W. Sanchirico, 'Inequality and uncertainty: Theory and legal applications', *U. Pa. L. Rev.*, **155**, 279, (2006).
- [2] Y. Bachrach, N. Betzler, and P. Faliszewski, 'Probabilistic possible winner determination', in *Proceedings of AAAI 2010*, (2010).
- [3] N. Bansal and M. Sviridenko, 'The santa claus problem', in *Proceedings of STOC'06*, pp. 31–40. ACM, (2006).
- [4] J. L. Bresina, 'Heuristic-Biased Stochastic Sampling', in *Proceedings of AAAI-96*, pp. 271–278, Portland, Oregon, (1996).
- [5] J. Broome, 'Uncertainty and fairness', *The Economic Journal*, **94**(375), 624–632, (1984).
- [6] E. Budish, 'The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes', *Journal of Political Economy*, **119**(6), 1061–1103, (2011).
- [7] 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).
- [8] H. Dalton, 'The measurement of the inequality of incomes', *The Economic Journal*, **30**(119), 348–361, (1920).
- [9] M. Fleurbaey, 'Assessing risky social situations', *The Journal of Political Economy*, **118**(4), 649–680, (2010).
- [10] T. Gajdos and J.-M. Tallon, 'Fairness under uncertainty', *Economics Bulletin*, **4**(18), 1–7, (2002).
- [11] J. C. Harsanyi, 'Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility', *Journal of Political Economy*, **63**, 309–321, (1955).
- [12] A. Holland and B. O'Sullivan, 'Robust solutions for combinatorial auctions', in *Proceedings of EC'2005*, pp. 183–192. ACM, (2005).
- [13] M. Lemaître, G. Verfaillie, and N. Bataille, 'Exploiting a common property resource under a fairness constraint: a case study', in *Proceedings of IJCAI-99*, pp. 206–211, (1999).
- [14] T. Lu and C. Boutilier, 'The unavailable candidate model : A decision-theoretic view of social choice', in *Proceedings of EC'10*, (2010).
- [15] J.-L. Marichal, *Aggregation Operators for Multicriteria Decision Aid*, Ph.D. dissertation, Faculty of Sciences of University of Liège, 1999.
- [16] A.W. Marshall, I. Olkin, and B.C. Arnold, *Inequalities: Theory of Majorization And Its Applications*, Springer Series in Statistics, Springer, 2nd edn., 2011.
- [17] H. Moulin, *Axioms of Cooperative Decision Making*, Cambridge University Press, 1988.
- [18] R. B. Myerson, 'Utilitarianism, egalitarianism, and the timing effect in social choice problems', *Econometrica*, **49**(4), 883–897, (1981).
- [19] A.C. Pigou, *Wealth and Welfare*, Macmillan, London, 1912.
- [20] H. Steinhaus, 'The problem of fair division', *Econometrica*, **16**(1), 101–104, (1948).
- [21] R. R. Yager, 'On ordered weighted averaging aggregation operators in multicriteria decision making', *IEEE Transactions on Systems, Man, and Cybernetics*, **18**, 183–190, (1988).